

Backup Exec 25.1 Administrator's Guide

Documentation version: Backup Exec 25.1

Legal Notice

Copyright ©2025 Arctera US LLC. All rights reserved.

Arctera and the Arctera Logo are trademarks or registered trademarks of Arctera US LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners. This product may contain third-party software for which Arctera is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Arctera product or available at:

<https://www.arctera.io/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and de-compilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Arctera US LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. ARCTERA US LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq." Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Arctera as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Arctera US LLC | www.arctera.io

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about our support offerings, you can visit our website at the following URL:

www.arctera.io/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.arctera.io/support

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Technical Support
 - Recent software configuration changes and network changes

Licensing and registration

If your product requires registration or a license key, access our technical support Web page at the following URL:

www.arctera.io/support

Customer service

Customer service information is available at the following URL:

www.arctera.io/support

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Advice about technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact us regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Worldwide

CustomerCare@arctera.io

Contents

Chapter 1	Introducing Backup Exec	35
	About Backup Exec	35
	How Backup Exec works	35
Chapter 2	Installation	38
	About the Backup Exec installation process and licensing options	39
	Backup Exec preinstallation checklist	41
	Running the Environment Check before installing or upgrading Backup Exec	42
	Microsoft SQL Server Express Edition components installed with Backup Exec	44
	System requirements for Backup Exec	45
	Installing Backup Exec by using the Installation Wizard	48
	Installing additional agents and features to the local Backup Exec server	58
	Push-installing Backup Exec to remote computers	60
	Methods for installing the Agent for Windows	68
	Push-installing the Agent for Windows to remote computers	68
	Installing updates to the Agent for Windows on remote computers	74
	Installing the Agent for Windows in an Active Directory network	75
	Using a command prompt to install the Agent for Windows on a remote computer	79
	Using a command script to install the Agent for Windows	83
	Installing the Remote Administrator	84
	Installing the Remote Administrator using the command line	86
	Running the Remote Administrator	88
	Installing Backup Exec using the command line (silent mode)	90
	Command line switches for silent mode installation of Backup Exec	91
	Creating and using installation parameter files	97

About the installation log	98
Viewing the Installation Summary Report	99
Repairing Backup Exec	99
About updates to Backup Exec using Arctera Update	100
Updating Backup Exec with Arctera Update	101
Viewing installed updates	103
Uninstalling Backup Exec updates	103
Viewing license information	104
Backup Exec license contract information	104
Updating expired license contracts	105
Managing license contract customer numbers	105
About upgrading to Backup Exec	106
Upgrade checklist for Backup Exec	108
Using the Migration Report to determine how existing jobs changed after an upgrade from a previous version of Backup Exec	109
Post-installation tasks	109
Uninstalling Backup Exec	110
Uninstalling Backup Exec using the command line	111
Uninstalling Backup Exec features from the local Backup Exec server	111

Chapter 3 Getting Started 113

About the Backup Exec Administration Console	113
Displaying the version information for Backup Exec	116
Locking and unlocking the Backup Exec Console	116
How to sort, filter, and copy information on the Backup Exec Administration Console	117
Customizing views on the Backup Exec Administration Console	119
Configuring the Home tab	120
Managing instance-based usage	125
Managing capacity-based usage	127
Managing restore only license	129
Configuring the RSS Reader	130
Creating a disaster preparation plan (DPP)	131
Getting started with backups	133

Chapter 4 Backups 136

How to prepare for your first backup with Backup Exec	138
Recommendations for when to use virtual-based backup and agent-based backup	139

Differences between the traditional and the forever incremental backups of virtual machines	141
Improving backup performance in Backup Exec	142
Troubleshooting backup performance	145
Required user rights for backup jobs	150
About the list of servers on the Backup and Restore tab	150
Adding servers that you want to back up to the list of servers on the Backup and Restore tab	151
Removing servers from the list of servers on the Backup and Restore tab	152
Creating a server group	153
Hiding or viewing server groups on the Backup and Restore tab	154
Adding servers to a server group	154
Removing servers from a server group	155
Editing a server group	155
Moving servers to the Retired Servers server group	156
Moving retired servers back to the All Servers server group	157
Removing server groups from Backup Exec	157
Backing up data	158
How Backup Exec backs up and restores volumes that are enabled with bit-locker	169
How Backup Exec backs up and restores volumes that are enabled for deduplication in Windows	170
How Backup Exec backs up and restores Microsoft Virtual Hard Disk (vhd) files	170
About selecting data to back up	171
Changing the order in which backup sources are backed up	176
Excluding files from backups	179
Including specific files with a backup job's backup selections	182
About backing up critical system components	184
Backup Exec Shadow Copy Components file system	188
Backup methods in Backup Exec	189
Configuring backup methods for backup jobs	197
How Backup Exec determines if a file has been backed up	198
Configuring Backup Exec to automatically delete files after a backup	202
Configuring network options for backup jobs	203
Running the next scheduled backup job before its scheduled time	204
Editing backup definitions	205
Testing credentials for backup sources	212
Replacing the credentials for a backup source	213

Creating new credentials for a backup source	214
Deleting retired or unused backup sources from the Credentials pane	215
How job scheduling works in Backup Exec	216
Including a specific date in the schedule for a backup job	217
Preventing backup jobs from running on a specific date	218
Viewing all scheduled backup jobs on a calendar	220
Adding a stage to a backup definition	220
Editing a stage	221
Duplicating backup sets or a job history manually	222
Running a test run job manually	226
Verifying backed up data manually	227
Copying data from a virtual tape library to a physical tape device using DirectCopy to tape	229
About using a self-protect backup job to protect server metadata	230

Chapter 5

Restores	234
Methods for restoring data in Backup Exec	235
Searching for data to restore	237
Configuring malware scan for virtual machine backup sets before restore	238
Restoring data from a server, a backup set, a backup job, or a storage device	241
Restoring file system data	244
Performing a complete online restore of a Microsoft Windows computer	246
Restoring System State	247
Installing a new Windows Server domain controller into an existing domain by using a redirected restore	249
Restoring Backup Exec Shadow Copy Components	251
Restoring utility partitions or Unified Extensible Firmware Interface system partitions	252
About restoring encrypted data	253
About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec	253
Canceling a restore job	253
How Backup Exec catalogs work	254
Configuring default options for catalogs	255
Moving the Backup Exec catalogs to a new directory	258
Cataloging backup sets	260

Chapter 6	Job management and monitoring	261
	How to monitor and manage jobs in Backup Exec	261
	About the Job Monitor	263
	Viewing job activity details for active jobs	265
	Canceling an active job	265
	Holding jobs	266
	Removing the hold on jobs	267
	Holding the job queue	268
	Removing the hold on the job queue	269
	Running a scheduled job immediately	270
	Changing the priority for a scheduled job	270
	Deleting scheduled jobs	272
	Retrying only failed resources	272
	Run backup job with debugging enabled	273
	About the Job History	274
	Viewing the history of a job	275
	Deleting a job from the Job History	275
	Running a job from the Job History	276
	Retrying only failed resources from the Job History	277
	Run backup job with debugging enabled from the Job History	278
	Viewing the job log	279
	Finding text in the job log	279
	Printing the job log	281
	Saving the job log	281
	Linking from the job log to the Arctera Technical Support website	282
	How to use job logs with vertical applications	282
	Configuring default job log options	283
	Error-handling rules for failed or canceled jobs	285
	Creating a custom error-handling rule	286
	Enabling or disabling error-handling rules	287
	Deleting a custom error-handling rule	287
	Enabling an error-handling rule for a failed job	287
	Custom error-handling rule for recovered jobs	288
	List of job statuses in Backup Exec	289
	Setting job status and recovery options	295
	About Anomaly Detection	297
	Managing anomaly detection	298

Chapter 7	Alerts and notifications	302
	Alerts and notifications in Backup Exec	303
	Where to find alerts in Backup Exec	304
	Enabling active alerts and alert history to display on the Home tab	306
	Active alerts properties	307
	Viewing the alert history for a server or a storage device	308
	Deleting an alert from alert history	308
	Copying alert text to a document or email	309
	Filtering alerts	309
	Viewing the job log from an alert	310
	Responding to active alerts	310
	Clearing all informational alerts manually	311
	Setting up notification for alerts	311
	Configuring email or text message notification for alerts	312
	Adding recipients for notification	314
	Adding a recipient group for alert notifications	316
	Removing a recipient from a group	317
	Disabling email or text message alert notification for a recipient	317
	Editing recipient notification properties	318
	Deleting recipients	318
	Configuring alert categories	319
	Assigning recipients to receive notifications for specific alert categories	321
	Sending a notification when a job completes	321
	Notification options for jobs	322
	Disabling notifications for a specific alert category	322
	Configuring default alert settings	323
	Enabling or disabling pop-up alerts	324
	SNMP traps for Backup Exec alerts	325
	Installing and configuring the SNMP system service	328
	Installing the Windows Management Instrumentation performance counter provider	328
	Installing the Windows Management Instrumentation provider for SNMP	328
	Uninstalling the Windows Management Instrumentation performance counter provider	329
	Uninstalling the Windows Management Instrumentation provider for SNMP	329

Chapter 8	Disk-based and network-based storage	330
	Features and types of disk-based storage and network-based storage	330
	Storage trending statuses for disk storage and virtual disks	332
	Setting low disk space thresholds on disk-based storage	333
	Configuring disk storage	334
	Changing the location of a disk storage device	337
	Editing disk storage properties	338
	How to restore data from a reattached or reinserted disk-based storage device	344
	Configuring disk cartridge storage	345
	Editing disk cartridge properties	346
	Editing disk cartridge media properties	350
	How data lifecycle management (DLM) deletes expired backup sets on disk-based storage	352
	Backup sets	359
	Changing the expiration date of backup sets on disk-based storage	362
	Retaining backup sets on disk-based storage to prevent them from expiring	363
	Releasing retained backup sets on disk-based storage	364
	Viewing the contents or properties of backup sets	365
Chapter 9	Cloud-based storage devices	367
	About cloud-based storage devices	368
	Amazon S3 cloud-based storage	368
	Requirements for configuring an Amazon S3 cloud-based storage device	369
	Configuring storage for Amazon cloud storage	369
	Google cloud-based storage	372
	Requirements for configuring a Google cloud-based storage device	372
	Configuring storage for Google cloud storage	373
	Microsoft Azure cloud-based storage	375
	Requirements for configuring a Microsoft Azure cloud-based storage device	375
	Configuring storage for Microsoft Azure cloud storage	376
	Private cloud-based storage	378
	Requirements for configuring a private cloud-based storage device	378
	Creating a cloud instance for a private cloud	379

Configuring storage for a private cloud-based storage device	380
Viewing and editing existing cloud instances for a private cloud	382
Deleting a cloud instance for a private cloud	382
About S3-Compatible Cloud Storage	383
Configuring S3-Compatible Cloud Storage with Backup Exec	384
Configuring S3-Compatible Cloud Storage with V4 authentication in Backup Exec	385
About Backup Exec Cloud Deduplication	388
Notes for Backup Exec Cloud Deduplication	389
Cloud deduplication storage device	390
Requirements for configuring a cloud deduplication storage device	390
About cloud immutable (WORM) storage support	393
Using the cloud admin command-line tool with Amazon S3	394
About immutable storage support for Azure blob storage	395
Using the cloud admin command-line tool with Azure blob storage	395
Configuring a cloud deduplication storage device	398
Deleting a cloud deduplication storage device	400
Notes for cloud-based storage devices	401
Editing the properties of a cloud-based storage device	402
Best practices for using cloud-based storage	406
Changing default cloud storage settings	407
About the Backup Exec™ CloudConnect Optimizer	408
Configuring the Backup Exec™ CloudConnect Optimizer	409

Chapter 10	Generic S3 Configurator	412
	About the Generic S3 Configurator	412
	Prerequisites for using Generic S3 Configurator	412
	Notes for Generic S3 Configurator	413
	Creating a cloud instance	413
	Deleting a cloud instance	415
	Adding a cloud region	415
	Viewing the cloud regions	416
	Updating a cloud region	417
	Deleting a cloud region	417

Chapter 11	OpenStorage devices	418
	Prerequisites for configuring OpenStorage devices	418
	Configuring an OpenStorage device	419
	Notes for OpenStorage devices	421
	Editing the properties of an OpenStorage device	422
	Data Lifecycle Management for WORM enabled OpenStorage devices	425
	Specifying a Backup Exec server that has proximity to a shared OpenStorage device	426
Chapter 12	Microsoft 365	428
	About support for Microsoft 365	429
	Requirements for Microsoft 365	429
	Configuring a tenant for Microsoft 365	430
	Backing up Microsoft 365 tenant data	433
	Supported workloads and entities for Microsoft 365	437
	Restoring Microsoft 365 tenant data	439
	Catalog operations for Microsoft 365	447
	CAS-MBES scenarios in Microsoft 365	447
	Notes for Microsoft 365	448
	Limitations of Microsoft 365	451
	Recommendations for Microsoft 365	458
	Best Practices for Microsoft 365	459
	OneDrive plugin: Performance and throttling configuration	460
	Tuning Scenarios	460
	Exchange plugin: Performance and throttling configuration	462
	Considerations before running a backup job	463
	Tuning Scenarios	463
	SharePoint plugin: Performance and throttling configuration	467
	Considerations before running a backup job	467
	Tuning Scenarios	468
	Teams plugin: Performance and throttling configuration	469
	Tuning scenario	470
Chapter 13	Tape storage	471
	Support for tape drives and robotic libraries	472
	Adding or replacing devices with the Hot-swappable Device Wizard	472
	Editing tape drive properties	473
	Viewing tape drive statistics	478
	Robotic libraries in Backup Exec	478

Requirements for setting up robotic library hardware	479
Inventorying robotic libraries when Backup Exec services start	480
Configuring barcode rules for a robotic library	480
Initializing a robotic library when the Backup Exec service starts	484
Defining a cleaning slot	485
Editing robotic library properties	485
Creating robotic library partitions	486
Adding or removing tape drives in a robotic library partition	488
Reassigning a slot base number for robotic libraries	488
Editing robotic library slot properties	488
Removing or reconfiguring robotic library partitions	489
Managing tapes	491
Default media sets	491
Overwrite protection periods and append periods in media sets	495
Creating media sets for tapes	499
Changing the overwrite protection period or the append period for a media set	502
Changing the name and description of a media set	504
Changing the media vault or vaulting periods for a media set	504
Deleting a media set	505
Renaming a media set	505
Media overwrite protection levels for tape media	505
Overwriting allocated or imported tape media	506
How Backup Exec searches for overwritable media in tape drives	506
Viewing tapes that are used by a job	509
Labeling tape media	509
About labeling imported tape media	511
How barcode labels become media labels for tapes	511
Renaming a tape media label	512
How WORM media is used in Backup Exec	512
Default media vaults	513
Changing the name or description of a media vault	514
Creating media vault rules to move tape media to and from media vaults	515
Updating the tape media location in media vaults	516
Deleting a media vault	516
Moving tape media to a vault	517

	Retiring damaged tape media	517
	Deleting tape media	518
	Erasing tape or disk cartridge media	518
	About cataloging tape or disk cartridge media that contains encrypted backup sets	520
	Associating tape media with a media set	520
	Editing tape media properties	521
	Tape media rotation strategies	523
Chapter 14	Storage device pools	527
	Creating storage device pools	527
	Specifying a default selection method for a device in a storage device pool	530
	Viewing jobs, job histories, and active alerts for a storage device pool	530
	Adding or removing devices in a storage device pool	531
Chapter 15	Storage operations	533
	About storage operation jobs	534
	Storage tab overview in Backup Exec	535
	Using the Configure Storage wizard	538
	Viewing details for multiple storage devices	540
	Sending a notification when a scheduled storage operation job completes	541
	Scheduling a storage operation job	541
	Editing global settings for storage	548
	Sharing storage devices	557
	Deleting a storage device	558
	Changing the state of a storage device to online	559
	Renaming a storage device	559
	Viewing jobs, job histories, backup sets, and active alerts for storage devices	560
	Cataloging a storage device	560
	Scanning a storage device	561
	Inventorying a storage device	562
	Inventorying and cataloging a storage device	563
	Pausing and unpausing a storage device	564
	Disabling and enabling a storage device	564
	Initializing a robotic library	564
	Formatting a tape as a WORM tape	565
	Retensioning a tape	565
	Formatting a tape in a tape drive	566

	Ejecting media from a disk cartridge or tape drive	567
	Cleaning a robotic library drive	567
	Importing media to Backup Exec	568
	Exporting media and expired media	576
	Locking and unlocking the robotic library's front portal	578
	Backup Exec server and storage device states	584
	Viewing the disk storage lockdown status	587
Chapter 16	Conversion to virtual machines	589
	How conversion of physical computers to virtual machines works	
	in Backup Exec	589
	Requirements for conversion to virtual machine jobs	597
	Converting to a virtual machine simultaneously with a backup job	
		599
	Converting to a virtual machine after a backup job	605
	Adding a conversion to virtual machine stage to a backup job	613
	Converting to a virtual machine from a point in time	620
	Creating a one-time conversion to a virtual machine	622
	Setting default options for conversion to virtual machine jobs	628
Chapter 17	Configuration and settings	632
	Configuring role-based access settings	634
	RBAC role details	637
	Changing default backup job settings	640
	Configuring schedules for backup jobs	644
	Configuring storage options for backup jobs	652
	Configuring automatic test run jobs for backup jobs	659
	Configuring automatic verify operations for backup jobs	661
	Configuring Instant GRT and full catalog options to improve backup	
	performance for GRT-enabled jobs	662
	Configuring Advanced Open File options for backup jobs	669
	Configuring checkpoint restart	672
	Configuring pre/post commands for backup or restore jobs	674
	Configuring parallel streams and job settings for Microsoft 365	679
	About preferred number of streams to use for backup	682
	Configuring file and folder options for backup jobs	683
	Setting default schedule options for rule-based jobs and run now	
	jobs	692
	Excluding dates from the backup schedule for all backups	695
	Removing dates from the list of excluded dates	697
	Exporting a list of dates that are excluded from all backups to another	
	server	698

Changing the default preferences	698
Configuring the default setting for backing up multiple servers or applications	700
Configuring database maintenance and security	701
Exporting the Backup Exec Database encryption key	705
Refreshing Backup Exec Database encryption keys	707
Configuring encryption for the connection to the Backup Exec Database	708
Scheduling Backup Exec to check logon accounts	712
Configuring Backup Exec to discover data to back up	713
Adding discovered servers to the list of servers in Backup Exec	715
Backup networks	716
Changing network and security options for Backup Exec	718
Disabling disk storage lockdown	724
Using Backup Exec with firewalls	725
Backup Exec ports	726
Backup Exec listening ports	728
About enabling a SQL instance behind a firewall	729
Using encryption with Backup Exec	729
Encryption key management	733
Creating encryption keys	734
Replacing an encryption key	735
Deleting encryption keys	738
Encryption keys and Salt	739
Granular Recovery Technology	739
Setting default Granular Recovery Technology (GRT) options	745
DBA-initiated job templates	746
Creating DBA-initiated job templates	747
Editing DBA-initiated job templates	748
Deleting DBA-initiated job templates	748
Storage options for DBA-initiated jobs	749
General options for DBA-initiated jobs	754
Network options for DBA-initiated jobs	754
Duplicate job settings for DBA-initiated jobs	755
Backup Exec logon accounts	759
Creating a Backup Exec logon account	762
Editing a Backup Exec logon account	763
Changing the password for a Backup Exec logon account	765
Replacing a Backup Exec logon account	766
Deleting a Backup Exec logon account	766
Changing your default Backup Exec logon account	768
Creating a new Backup Exec System Logon Account	768

Copying logon account information to another Backup Exec server	769
Testing logon accounts	770
Starting and stopping Backup Exec services	771
Changing the credentials for a service account	771
Changing startup options for Backup Exec services	773
Configuring audit logs	774
Viewing the audit log	775
Removing entries from the audit log	775
Saving an audit log to a text file	776
Copying configuration settings to another Backup Exec server	776
Viewing server properties	777
Configuring default backup settings	778
Changing malware scan options	783

Chapter 18

Reports	785
Reports in Backup Exec	785
Running a report now	787
Scheduling a report	787
Creating a custom report	789
Adding or removing fields on a custom report	794
Changing filters for a custom report	794
Changing the way data is grouped or sorted in a custom report	796
Changing graph options in custom reports	797
Previewing custom reports	798
Copying a custom report	798
Saving a report	798
Printing a report from the Backup Exec Report Viewer	799
Viewing completed reports	800
Editing a report	800
Re-running a completed report	800
Deleting a report	801
Setting defaults for standard and custom reports	801
Viewing report properties	802
List of Backup Exec standard reports	803
Alert History report	808
Alert History By Backup Exec Server report	808
Anomaly Detection Summary report	809
Audit Log report	810
Backup Job Success Rate report	811
Backup Recommendations report	811

Backup Resource Success Rate report	812
Backup Sets by Media Set report	812
Backup Size By Resource report	813
Cloud Storage Summary report	813
Daily Device Utilization report	814
Deduplication Disk and Cloud Deduplication Device Summary report	815
Deduplication Summary report	816
Device Summary report	817
Disk Storage Summary report	817
Error-Handling Rules report	818
Event Recipients report	819
Failed Backup Jobs report	820
Jobs Summary report	821
Malware Scan Summary report	821
Managed Backup Exec Servers report	822
Media Audit report	824
Media Errors report	825
Media Required for Recovery report	825
Media Summary report	826
Media Vault Contents report	827
Move Media to Vault report	827
Operations Overview report	828
Overnight Summary report	830
Problem Files report	831
Recently Written Media report	832
Recovery Ready Validation Summary	832
Resource Protected Recently report	833
Resource Risk Assessment report	834
Restore Set Details by Resource report	835
Retrieve Media from Vault report	836
Robotic Library Inventory report	836
Scheduled Server Workload report	837
Scratch Media Availability report	838
Test Run Results report	839
User Role Details report	839

Chapter 19	GDPR Guard	841
	About GDPR Guard	841
	Backup Exec Management Command Line (BEMCLI) commands for import and export	843
	Supported types of backed up data	845

	How to block access to backed up items	846
	Restoring blocked items	847
	Best practices for blocking access to backed up items with GDPR Guard	848
Chapter 20	Troubleshooting Backup Exec	850
	Troubleshooting hardware-related issues in Backup Exec	851
	Troubleshooting robotic libraries and tape drives	853
	How to get more information about alerts and error messages	857
	Troubleshooting backup issues in Backup Exec	858
	Troubleshooting failed components in the SAN	860
	Troubleshooting offline storage devices in a SAN	860
	Finding hardware errors in a SAN	862
	Resetting the SAN	863
	Bringing storage devices online after an unsafe device removal event in a SAN	864
	Troubleshooting installation issues in Backup Exec	864
	Troubleshooting blocked access to backed up items with GDPR Guard	865
	How to improve Backup Exec's performance	867
	Accessing Arctera Online	867
	Searching the Arctera Knowledge Base	868
	Contacting Backup Exec Technical Support	869
	Using Remote Assistance	869
	Managing your Backup Exec support cases	869
	About Backup Exec diagnostic tools	869
	Running the Veritas QuickAssist Help Tool	870
	Generating a diagnostic file for troubleshooting Backup Exec	871
	Command line switches for a diagnostic file	872
	Running the begather utility to troubleshoot Backup Exec components on Linux servers	875
	Using the Backup Exec Debug Monitor for troubleshooting	875
	About the Backup Exec debug tool	876
Chapter 21	Simplified Disaster Recovery	877
	About Simplified Disaster Recovery	877
	Requirements for using Simplified Disaster Recovery	878
	Preparing computers for use with Simplified Disaster Recovery	881
	How to ensure that backups are enabled for Simplified Disaster Recovery	885
	How Simplified Disaster Recovery uses disaster recovery information files	887

Setting or changing the alternate location for the disaster recovery information file	889
Changing the default path for the disaster recovery information files	890
Disaster recovery information file data paths	890
Creating a Simplified Disaster Recovery disk image	892
Contents of the Simplified Disaster Recovery disk image	909
User scenarios when a user starts the Create Recovery Disk Wizard	910
Preparing to recover from a disaster by using Simplified Disaster Recovery	913
Hardware replacement during disaster recovery	914
Prepare to recover IBM computers with Simplified Disaster Recovery	915
Recovering a computer with Simplified Disaster Recovery	916
Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces	922
Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature	924
Recovery notes for using Simplified Disaster Recovery with Windows BIOS system	925
Advanced Disk Configuration on the Recover This Computer Wizard	926
Performing manual disaster recovery	929
Performing manual disaster recovery of a local Backup Exec server on a Windows computer	930
Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer	935

Chapter 22	Forever Incremental Backup	940
	About Forever Incremental Backup	940
	How do forever incremental backups work	941
	Supported storage in forever incremental backups	946
	Backing up virtual machines using forever incremental backups	947
	CAS-MBES scenarios in forever incremental backups	948
	Notes for forever incremental backups	948
	Recommendations for forever incremental backups	949
	Limitations of forever incremental backups	950

Appendix A	Backup Exec Agent for Windows	951
	About the Agent for Windows	951
	Requirements for the Agent for Windows	952
	Stopping and starting the Agent for Windows	953
	Establishing a trust between the Backup Exec server and a remote computer	953
	About the Backup Exec Agent Utility for Windows	954
	Starting the Backup Exec Agent Utility	955
	Viewing the activity status of the remote computer in the Backup Exec Agent Utility	955
	Viewing the activity status of the remote computer from the system tray	956
	Starting the Backup Exec Agent Utility automatically on the remote computer	957
	Setting the refresh interval on the remote computer	957
	About publishing the Agent for Windows to Backup Exec servers	958
	Adding Backup Exec servers that the Agent for Windows can publish to	958
	Editing Backup Exec server information that the Agent for Windows publishes to	959
	Removing Backup Exec servers that the Agent for Windows can publish to	960
	Configuring database access for Oracle operations	960
	Removing a security certificate for a Backup Exec server that has a trust with the Agent for Windows	963
	Using the Backup Exec Agent Utility Command Line Applet	963
	Backup Exec Agent Utility Command Line Applet switches	964
Appendix B	Backup Exec Deduplication Feature	968
	About the Deduplication feature	969
	Deduplication methods for Backup Exec agents	971
	Requirements for the Deduplication feature	972
	Configuring a 125 TB Deduplication storage folder in Backup Exec	977
	Installing the Deduplication feature	980
	Converting an older version of Deduplication Storage to a newer version	980
	Creating or importing deduplication disk storage	982
	Editing the properties of a deduplication disk storage device	985

Changing the password for the logon account for deduplication disk storage	990
Selecting storage devices for direct access sharing	991
Editing server properties for direct access	992
Changing the location of a deduplication disk storage device	993
Sharing a deduplication device between multiple Backup Exec servers	995
How to use client-side deduplication	995
How to set up backup jobs for deduplication	996
Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices	997
Copying deduplicated data to tapes	1000
Using deduplication with encryption	1000
Restoring a deduplication disk storage device or deduplicated data	1001
Disaster recovery of deduplication disk storage devices	1001
Disaster recovery of OpenStorage devices	1002

Appendix C	Backup Exec Agent for VMware	1004
	About the Agent for VMware	1005
	Requirements for using the Agent for VMware	1005
	Using the Agent for VMware with Windows Server 2016 or later	1007
	About installing the Agent for VMware	1008
	Adding VMware vCenter Servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab	1008
	Viewing details about VMware resources	1009
	Installing the Agent for Windows on VMware virtual machines	1010
	Push-installing the Agent for Windows to VMware virtual machines	1011
	About establishing trust for a vCenter/ESX(i) server	1012
	Establishing trust for a vCenter/ESX(i) server	1013
	Backing up VMware virtual machines	1013
	Setting default backup options for virtual machines	1022
	How Backup Exec automatically backs up new VMware virtual machines during a backup job	1027
	Using Granular Recovery Technology (GRT) with the Agent for VMware	1028
	How cataloging works with VMware virtual machine backups	1031
	Multi-stream backup of virtual machines	1033

Restoring VMware virtual machines and vmdk files	1034
About instant recovery of a VMware virtual machine	1044
Requirements for instant recovery of a VMware virtual machine	1047
Notes about instant recovery of a VMware virtual machine	1048
Best Practices for instant recovery of a VMware virtual machine	1049
Creating an instant recovery job for a VMware virtual machine	1049
Removing an instantly recovered VMware virtual machine	1052
Troubleshooting the Agent for VMware	1053
About Recovery Ready for VMware virtual machines	1054
Requirements for validating a VMware virtual machine for recovery	1057
Notes about validating a VMware virtual machine for recovery	1058
Best Practices for validating VMware virtual machine for recovery	1059
Creating a validate virtual machine for recovery job	1059

Appendix D	Backup Exec Agent for Microsoft Hyper-V	1063
	About the Agent for Microsoft Hyper-V	1064
	Requirements for using the Agent for Microsoft Hyper-V	1066
	About installing the Agent for Microsoft Hyper-V	1068
	Notes about using the Agent for Hyper-V	1068
	Disk space optimization with the Agent for Hyper-V	1070
	Adding a Hyper-V host to the list of servers on the Backup and Restore tab	1070
	Viewing details about Hyper-V resources	1071
	Installing the Agent for Windows on Hyper-V virtual machines	1072
	Push-installing the Agent for Windows to Hyper-V virtual machines	1073
	Backing up Microsoft Hyper-V virtual machines	1074
	Setting default backup options for Hyper-V	1083
	How Backup Exec automatically protects new virtual machines during a backup job	1086
	Using Granular Recovery Technology (GRT) with the Agent for Hyper-V	1087
	About backing up and restoring highly available Hyper-V virtual machines	1090

About backing up VMs hosted on SMB/Scale out File Server	1091
How cataloging works with Hyper-V virtual machine backups	1091
Restoring Microsoft Hyper-V virtual machines	1093
About instant recovery of a Hyper-V virtual machine	1098
Requirements for instant recovery of a Hyper-V virtual machine	1103
Notes about instant recovery of a Hyper-V virtual machine	1103
Creating an instant recovery job for a Hyper-V virtual machine	1104
Post-instant recovery tasks	1106
About removing an instantly recovered Hyper-V virtual machine	1107
Removing an instantly recovered Hyper-V virtual machine	1108
Best practices for instant recovery of a Hyper-V virtual machine	1109
About Recovery Ready for Hyper-V virtual machines	1110
Requirements for validating a Hyper-V virtual machine for recovery	1114
Notes about validating a Hyper-V virtual machine for recovery	1114
Best practices about validating a Hyper-V virtual machine for recovery	1115
Creating a validate virtual machine for recovery	1115
Troubleshooting issues with Backup Exec Agent for Microsoft Hyper-V	1117

Appendix E	Backup Exec Agent for Microsoft SQL Server	1119
	About the Agent for Microsoft SQL Server	1119
	Requirements for using the SQL Agent	1121
	About installing the SQL Agent	1121
	Backup strategies for SQL	1121
	Adding SQL Servers to the list of servers on the Backup and Restore tab	1123
	Configuring Backup Exec to run a consistency check before every SQL backup	1124
	Using snapshot technology with the SQL Agent	1125
	Using database snapshots for SQL Server	1126
	Backing up SQL databases and transaction logs	1127

Setting default backup options for SQL Server	1138
Restoring SQL databases and transaction logs	1146
Restoring the SQL master database	1148
Restarting SQL using database copies	1149
Disaster recovery of a SQL Server	1151
Manual recovery of a SQL Server	1153
About SQL Server Always On availability groups	1154
Requirements of Always On availability groups from SQL Server	1154
Terms used by SQL Server Always On availability groups	1154
Backup Exec recommendations for SQL Server Always On availability groups	1156
Adding a listener to the list of servers on the Backup and Restore tab	1157
Backing up databases from a SQL Server availability group	1158
Restoring databases from a SQL Server availability group	1159

Appendix F	Backup Exec Agent for Microsoft Exchange Server	1160
	About the Backup Exec Exchange Agent	1161
	Requirements for using the Exchange Agent	1161
	Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations	1165
	About installing the Exchange Agent	1170
	Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab	1171
	Managing preferred server configurations for Microsoft Exchange Database Availability Groups	1171
	Recommended configurations for Exchange	1174
	Requirements for accessing Exchange mailboxes	1175
	Backup strategies for Exchange	1176
	How Granular Recovery Technology works with the Exchange Information Store	1179
	Snapshot and offhost backups with the Exchange Agent	1180
	Backing up Exchange data	1182
	Setting default backup options for Exchange Server	1190
	Restoring Exchange data	1196
	Disaster recovery of an Exchange Server	1202

Appendix G	Backup Exec Agent for Microsoft SharePoint	1204
	About the Agent for Microsoft SharePoint	1204
	About installing the Agent for Microsoft SharePoint	1205
	Requirements for the Agent for Microsoft SharePoint	1205
	Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation	1206
	Adding a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab	1207
	Backing up Microsoft SharePoint data	1207
	Setting default backup options for SharePoint	1211
	Restoring Microsoft SharePoint data	1214
	Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec	1216
	Viewing or changing SharePoint farm properties	1217
	Disaster recovery of Microsoft SharePoint data	1218
Appendix H	Backup Exec Agent for Oracle on Windows or Linux Servers	1223
	About the Backup Exec Oracle Agent	1223
	About installing the Oracle Agent	1225
	Configuring the Oracle Agent on Windows computers and Linux servers	1225
	Configuring an Oracle instance on Windows computers	1226
	Viewing an Oracle instance on Windows computers	1230
	Editing an Oracle instance on Windows computers	1231
	Deleting an Oracle instance on Windows computers	1232
	Enabling database access for Oracle operations on Windows computers	1232
	Configuring an Oracle instance on Linux servers	1235
	Viewing an Oracle instance on Linux servers	1237
	Editing an Oracle instance on Linux servers	1238
	Deleting an Oracle instance on Linux servers	1238
	Enabling database access for Oracle operations on Linux servers	1239
	About authentication credentials on the Backup Exec server	1240
	Setting authentication credentials on the Backup Exec server for Oracle operations	1241
	Deleting an Oracle server from the Backup Exec server's list of authentication credentials	1242
	About Oracle instance information changes	1243
	About backing up Oracle databases	1243

About backing up Oracle RAC databases	1245
About performing a DBA-initiated backup job for Oracle	1246
Oracle backup options	1247
About restoring Oracle resources	1248
About DBA-initiated restore for Oracle	1251
Oracle restore options	1252
About redirecting a restore of Oracle data	1252
Oracle advanced restore options	1253
Performing a redirect restore of an Oracle 12c and later supported database using Backup Exec	1254
Requirements for recovering the complete Oracle instance and database using the original Oracle server	1258
Recovering the complete Oracle instance and database using the original Oracle server	1258
Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server	1259
Recovering the complete Oracle instance or database to a computer other than the original Oracle server	1260
Best practices for Backup Exec Agent for Oracle on Windows and Linux Servers	1261

Appendix I	Backup Exec Agent for Enterprise Vault	1265
	About the Agent for Enterprise Vault	1265
	Requirements for the Enterprise Vault Agent	1269
	About installing the Enterprise Vault Agent	1270
	About backup methods for Enterprise Vault backup jobs	1271
	Enterprise Vault backup options	1274
	About backing up Enterprise Vault components	1274
	About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases	1277
	Restoring Enterprise Vault	1277
	Enterprise Vault restore options	1279
	About restoring individual files and folders with the Enterprise Vault Agent	1282
	About automatic redirection of Enterprise Vault components under an Enterprise Vault server	1282
	Redirecting a restore for an Enterprise Vault component	1283
	Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database	1285
	Best practices for the Enterprise Vault Agent	1287
	About the Backup Exec Migrator for Enterprise Vault	1287

Backup Exec Migrator for Enterprise Vault requirements	1288
How the Backup Exec Migrator works	1288
About using staged migrations with Backup Exec and the Backup Exec Migrator	1292
About Backup Exec Migrator events	1293
About Backup Exec Migrator logs	1293
How to enable Backup Exec Migrator logging	1294
About deleting files migrated by Backup Exec Migrator	1296
Configuring the Backup Exec Migrator	1297
About viewing migrated Enterprise Vault data	1304
About retrieving migrated Enterprise Vault data	1305
About the Partition Recovery Utility	1306
Enterprise Vault logon account	1309
Enterprise Vault options	1310

Appendix J Backup Exec Agent for Microsoft Active Directory

About the Agent for Microsoft Active Directory	1311
Requirements for the Agent for Microsoft Active Directory	1312
About backing up Active Directory and ADAM/AD LDS	1313
Editing options for Active Directory and ADAM/AD LDS backup jobs	1314
Microsoft Active Directory backup job options	1316
About restoring individual Active Directory and ADAM/AD LDS objects	1317

Appendix K Backup Exec Central Admin Server Feature

About the Central Admin Server feature	1321
Requirements for installing CAS	1323
How to choose the location for CAS storage and media data	1324
About installing the Central Admin Server feature	1326
Push-installing a managed Backup Exec server from the central administration server	1327
Managed Backup Exec Server Configuration options	1331
Installing a managed Backup Exec server across a firewall	1333
Upgrading an existing CAS installation	1335
Changing a Backup Exec server to a central administration server	1337
Changing a Backup Exec server to a managed Backup Exec server	1338
Deleting a managed Backup Exec server from a CAS environment	1339

Renaming a central administration server	1341
Renaming a managed Backup Exec server	1342
How to reduce network traffic in CAS	1343
CAS distributed, centralized, and replicated catalog locations	1344
Changing the settings for a managed Backup Exec server	1346
What happens when CAS communication thresholds are reached	1357
Enabling or disabling communications between the managed Backup Exec server and the central administration server	1358
Alerts and notifications in CAS	1358
Enabling managed Backup Exec servers to use any available network interface card	1360
About job delegation in CAS	1360
About copying jobs instead of delegating jobs in CAS	1361
About adding storage devices in a CAS environment	1361
How data lifecycle management (DLM) works in a CAS environment	1361
Obtaining media audit information for a managed Backup Exec server	1362
How to use Backup Exec server pools in CAS	1363
Selecting a Backup Exec server pool for backups	1364
Creating a Backup Exec server pool	1364
Adding managed Backup Exec servers to a Backup Exec server pool	1365
Deleting a Backup Exec server pool	1365
Removing a managed Backup Exec server from a Backup Exec server pool	1366
How centralized restore works in CAS	1366
How CAS restores data that resides on multiple storage devices	1367
About recovering failed jobs in CAS	1369
Pausing or resuming a managed Backup Exec server	1370
Stopping or starting Backup Exec services for a managed Backup Exec server	1371
Viewing managed Backup Exec server properties	1372
Viewing the settings for a central administration server	1375
Disaster recovery in CAS	1377
Troubleshooting CAS	1378
Running the Backup Exec Utility for CAS operations	1380
Uninstalling Backup Exec from the central administration server	1380
Uninstalling Backup Exec from a managed Backup Exec server	1381

Appendix L	Backup Exec Advanced Disk-based Backup Feature	1382
	About the Advanced Disk-based Backup feature	1382
	How to use synthetic backups in place of recurring full backups	1383
	Setting default backup options for the Advanced Disk-based Backup feature	1385
	About true image restore for synthetic backups	1386
	How to use off-host backup to process remote computer backups on the Backup Exec server	1388
	Configuring off-host backup options for a backup job	1391
	Best practices for off-host backup	1393
	Troubleshooting off-host backups	1394
	Off-host backup issues with hardware providers	1397
Appendix M	Backup Exec NDMP Feature	1398
	Features of the NDMP feature	1399
	Requirements for using the NDMP feature	1399
	About installing the NDMP feature	1400
	Adding NDMP servers to Backup Exec	1400
	Sharing the tape drives on NDMP servers between multiple Backup Exec servers	1402
	NDMP backup options for NDMP servers	1403
	NDMP server backup selections	1406
	How to use patterns to exclude files and directories from an NDMP server backup	1406
	Supported configurations for duplicating data from NDMP servers	1408
	About restoring and redirecting restore data for NDMP servers	1408
	NDMP server restore options	1409
	Setting the default backup options for the NDMP feature	1412
	Viewing the properties of an NDMP server	1413
	NDMP server properties	1414
	Viewing storage properties for an NDMP server	1414
	Storage properties for an NDMP server	1414
	About support for NetApp's Data ONTAP	1415
Appendix N	Backup Exec File Servers	1416
	About AWS FSx for Windows File Server	1416
	About Azure Files	1416
	Pre-requisites for AWS FSx and Azure Files	1417

Notes for AWS FSx and Azure Files	1417
Recommendation for AWS FSx and Azure Files	1417
Best practices for AWS FSx and Azure Files	1418
Adding AWS FSx or Azure Files to Backup Exec	1418
Backing up AWS FSx or Azure Files	1418
Restoring AWS FSx or Azure Files	1419
Limitation of Azure Files	1420
 Appendix O Backup Exec Agent for Linux and Unix	 1421
About the Agent for Linux and Unix	1422
About open files and the Agent for Linux	1422
Requirements for the Agent for Linux and Unix	1423
About installing the Agent for Linux and Unix	1423
Installing the Agent for Linux and Unix	1424
About the Backup Exec operators (beoper) group for the Agent for Linux and Unix	1427
About establishing trust for a remote Linux and Unix computer in the Backup Exec list of servers	1428
Establishing trust and adding a remote Linux and Unix computer to the Backup Exec list of servers	1429
Adding additional Backup Exec servers to which the Agent for Linux and Unix can publish information	1429
About configuring the Agent for Linux and Unix	1430
Excluding files and directories from all backup jobs for Linux and Unix computers	1431
Editing configuration options for Linux and Unix computers	1432
Configuration options for Linux and Unix computers	1432
About backing up a Linux and Unix computer by using the Agent for Linux and Unix	1439
Linux and Unix backup options	1439
About backing up Linux and Unix shares without using the Agent for Linux and Unix	1441
About restoring data to Linux and Unix computers	1441
Restore job options for Linux and Unix computers	1442
Editing the default backup job options for Linux and Unix computers	1443
Default backup job options for Linux and Unix computers	1443
Uninstalling the Agent for Linux and Unix	1445
Manually uninstalling the Agent for Linux and Unix	1446
Run-time scripts to remove when manually uninstalling the Agent for Linux and Unix	1447
Starting the Agent for Linux and Unix daemon	1449

Stopping the Agent for Linux and Unix daemon	1449
Troubleshooting the Agent for Linux and Unix	1450
Appendix P Backup Exec Remote Media Agent for Linux	1455
About the Remote Media Agent for Linux	1456
How the Remote Media Agent for Linux works	1457
Requirements for the Remote Media Agent for Linux	1458
About open files and the Remote Media Agent for Linux	1459
About installing the Remote Media Agent for Linux	1459
Installing the Remote Media Agent for Linux	1460
Uninstalling the Remote Media Agent for Linux	1463
Starting the Remote Media Agent for Linux daemon	1464
Stopping the Remote Media Agent for Linux daemon	1464
About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers	1465
Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers	1465
Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish	1466
Finding simulated tape library files	1466
About the Backup Exec operators (beoper) group for the Remote Media Agent for Linux	1467
Creating the Backup Exec operators (beoper) group manually for the Remote Media Agent for Linux	1468
About adding a Linux server as a Remote Media Agent for Linux	1468
Adding a Linux server as a Remote Media Agent for Linux	1469
Remote Media Agent for Linux options	1469
Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux	1471
About creating storage device pools for devices attached to the Remote Media Agent for Linux	1471
Editing properties for the Remote Media Agent for Linux	1472
Remote Media Agent for Linux properties	1472
Deleting a Remote Media Agent for Linux from the Backup Exec list of servers	1473
Sharing a Remote Media Agent for Linux between multiple Backup Exec servers	1473
About backing up data by using the Remote Media Agent for Linux	1474
About restoring data by using the Remote Media Agent for Linux	1474

About the Tape Library Simulator Utility	1475
Creating a simulated tape library	1476
Simulated Tape Library options	1477
Viewing simulated tape libraries properties	1477
Simulated tape library properties	1478
Deleting a simulated tape library	1479
Managing simulated tape libraries from the command line	1480
Command line switches for the Tape Library Simulator Utility	1480
Troubleshooting the Remote Media Agent for Linux	1481
Glossary	1484
Index	1490

Introducing Backup Exec

This chapter includes the following topics:

- [About Backup Exec](#)
- [How Backup Exec works](#)

About Backup Exec

Backup Exec is a high-performance data management solution for Windows® servers networks. With its client/server design, Backup Exec provides fast, reliable backup and restore capabilities for servers, applications, and workstations across the network.

Backup Exec is available in several configurations that can accommodate networks of all sizes. In addition, Backup Exec's family of agents and features offers solutions for scaling your Backup Exec environment and extending platform and feature support.

See [“How Backup Exec works”](#) on page 35.

How Backup Exec works

You use the Backup Exec Administration Console to interact with Backup Exec for tasks like submitting backups, restoring data, configuring storage, and monitoring jobs. You can run the Administration Console from the Backup Exec server, which is a Windows server on which Backup Exec is installed, or from a remote computer. After backups, restores, or other operations are created, the Backup Exec server processes the jobs or delegates the jobs for processing in multi-Backup Exec server environments.

Backup Exec includes the following features:

- Convenient backup scheduling

Backup Exec administrators can set up scheduled backups for Windows and Linux computers across the network. Backup Exec's flexible calendar-based administration lets you easily schedule backups for processing during off-peak hours.

- **Complete system recovery**
Backup Exec's Simplified Disaster Recovery takes all of the guesswork out of recovering an entire system. While configuring a backup, you get a clear indication that you have selected the data that is necessary to perform a Simplified Disaster Recovery-enabled backup. After you back up a computer's critical system components, use the **Create Simplified Disaster Recovery Disk Wizard** to create a Simplified Disaster Recovery disk image. You can then use the disk image to perform a disaster recovery of the computer.
- **Comprehensive monitoring and intuitive mechanisms for everyday tasks**
The **Job Monitor** provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The **Home** tab lets you view statistics for your entire Backup Exec environment. From the **Servers** view, you can monitor the backup status for all of the computers on your network. Interactive alerts display the situations that require your attention. Backup results can be viewed from a backup's job history. The job history contains statistics, errors, and other information pertaining to the backup. Backup Exec's catalog is a database of all backed-up data, and Backup Exec uses the catalog to track restore selections.
Wizards guide you through most Backup Exec operations, including the creation of backup and restore jobs, and the configuration of storage.
- **Automated data lifecycle management for disk-based and cloud storage**
Backup Exec uses data lifecycle management to automatically delete expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, cloud storage, and virtual disks. You specify how long to keep backup data when you create a backup job that is sent to a disk-based storage device. When the amount of time to keep the backup data expires, the data lifecycle management feature deletes the backup sets and reclaims the disk space for use by new backup sets.
[See "How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage" on page 352.](#)
- **Simplified device and media management**
Backup Exec uses the Advanced Device and Media Management (ADAMM) feature to manage data retention on tapes. ADAMM expires the backup sets that are stored on media according to a set of rules that you apply to the media.
- **Backup Exec on AWS and Azure Marketplace**

Backup Exec is available on the AWS and Azure Marketplace as a trial version, allowing you to evaluate the features, performance, scalability, and compatibility with your specific business needs. This approach minimizes the risk associated with adopting a new technology. By trying Backup Exec through the marketplace, you can get a hands-on product experience. There is additional information available about deploying Backup Exec on AWS and Azure Marketplace. For more information, refer to the Veritas Open eXchange documentation.

Note: You need to create a virtual machine with the appropriate resources and use the Backup Exec image published on the marketplace. Before you use Backup Exec, ensure that you run Windows updates.

Installation

This chapter includes the following topics:

- [About the Backup Exec installation process and licensing options](#)
- [Backup Exec preinstallation checklist](#)
- [Running the Environment Check before installing or upgrading Backup Exec](#)
- [Microsoft SQL Server Express Edition components installed with Backup Exec](#)
- [System requirements for Backup Exec](#)
- [Installing Backup Exec by using the Installation Wizard](#)
- [Installing additional agents and features to the local Backup Exec server](#)
- [Push-installing Backup Exec to remote computers](#)
- [Methods for installing the Agent for Windows](#)
- [Installing the Remote Administrator](#)
- [Installing Backup Exec using the command line \(silent mode\)](#)
- [About the installation log](#)
- [Viewing the Installation Summary Report](#)
- [Repairing Backup Exec](#)
- [About updates to Backup Exec using Arctera Update](#)
- [Viewing installed updates](#)
- [Uninstalling Backup Exec updates](#)
- [Viewing license information](#)

- [Backup Exec license contract information](#)
- [About upgrading to Backup Exec](#)
- [Post-installation tasks](#)
- [Uninstalling Backup Exec](#)
- [Uninstalling Backup Exec using the command line](#)
- [Uninstalling Backup Exec features from the local Backup Exec server](#)

About the Backup Exec installation process and licensing options

The Backup Exec installation wizard guides you through the process of installing Backup Exec and its agents and features. Using the installation wizard, you can install Backup Exec and its agents and features on a local computer or you can push-install them to a remote computer. The computer on which Backup Exec is installed is called the Backup Exec server. Additionally, you can install the Remote Administrator, which lets you administrate the Backup Exec server from a remote Windows server or workstation.

Note: You cannot install Backup Exec or the Agent for Windows on a volume that has been enabled for data deduplication in Windows, on an ReFS volume, or on Cluster Shared Volumes.

When you install Backup Exec, you can input a license in two ways:

Table 2-1 Licensing options

Item	Description
Enter entitlement IDs manually	You can enter the entitlement IDs that are listed on your license certificate. Entitlement IDs contain one letter and 10 numbers, such as A0123456789. After you add your entitlement IDs, enter the Veritas User Account credentials to connect to the Licensing Portal. Backup Exec connects to the Licensing Portal and downloads the license files. An Internet connection is required to enter entitlement IDs manually. Note: Licenses from previous versions of Backup Exec cannot activate the current version of Backup Exec.
Import licenses from the License File	You can import your License Files (.slf) from a network share or from a local drive. You must go to the Licensing Portal to obtain them. From the Licensing Portal, you receive one file with all of the entitlement IDs that you registered. After installation, the .slf files can be found in the following location: Windows Server 2012 and later: C:\Programdata\Veritas Shared\Licenses
Install a 60-day trial version	A 60-day trial version is available for Backup Exec. If you do not enter any entitlement IDs or license files during the installation process, a trial version is installed.

After Backup Exec has been installed, you can install additional agents and features if you have valid licenses for them. For more information about how to add agents and features, refer to the Backup Exec Price and Licensing Guide. You can also push-install Backup Exec or the Agent for Windows to remote computers.

Installation from a command line is also available. Command line installation is called silent mode installation. The silent mode installation uses the Setup.exe program on the Backup Exec installation media.

Backup Exec may install the following additional products during the installation process:

- Microsoft Report Viewer 2015
- Microsoft.NET Framework 4.8
- Microsoft Visual C++ 2012 Redistributable Package
- Microsoft Visual C++ 2015-2022 Redistributable Package
- Microsoft SQL Server Express

For information about the best practices to install Backup Exec, refer to *Backup Exec Best Practices*.

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 48.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“Push-installing Backup Exec to remote computers”](#) on page 60.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 68.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 90.

See [“Backup Exec preinstallation checklist”](#) on page 41.

Backup Exec preinstallation checklist

Before you install Backup Exec, you should do the following:

- Run the Backup Exec Environment Check on the computer on which you want to install Backup Exec. The Environment Check analyzes the computer to make sure that the installation process can complete. If Backup Exec finds any configuration issues that can be fixed during the installation, or that may prevent the installation, warnings appear. Although the Environment Check runs automatically during installation, you may want to run it manually before you install Backup Exec or before you back up data with Backup Exec.
See [“Running the Environment Check before installing or upgrading Backup Exec”](#) on page 42.
- Check the Backup Exec Hardware Compatibility List to confirm that your storage device hardware is supported in this version of Backup Exec.
- Install the storage device hardware (controller, drives, robotic libraries) on the Backup Exec server. Refer to the documentation that is included with your storage device hardware for installation instructions. Use the appropriate Windows hardware setup functions to configure your controller and storage devices. Refer to your Microsoft Windows documentation for more information. You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.

- Check your Windows security settings to make sure that they work properly with the Backup Exec service account.
See [“Changing the credentials for a service account”](#) on page 771.
- Ensure that port 50104 is available for use by the Backup Exec Management Service.
- If the drive on which you want to install Backup Exec is encrypted or compressed, and you want to use a default SQL Express database, verify that an unencrypted and uncompressed drive is available for SQL Express installation
- Check the computer name of the computer on which you want to install Backup Exec. It should only use standard ANSI characters. You may receive errors if you install Backup Exec on a computer with a name that uses non-standard characters.
- If you want to install Backup Exec to a non-English version of Windows, download the Microsoft SQL Server Express setup file for the language that you want to install from the Microsoft website before you install Backup Exec if all of the following are true:
 - You want to use a local Backup Exec SQL Express instance.
 - You have non-English SQL Server instances on the computer on which you want to install Backup Exec.

Running the Environment Check before installing or upgrading Backup Exec

The Backup Exec Environment Check is a utility that runs on a computer automatically during installation and that reports the following:

- If the computer meets the minimum requirements for installation, such as the operating system, disk and physical memory, and sufficient logon account privileges.
See [“System requirements for Backup Exec”](#) on page 45.
- If the third-party software that uses Backup Exec ports is configured correctly.
- If required components are installed, and if they are compatible with Backup Exec.
- If previous versions of Backup Exec and Backup Exec features are installed.
- If storage device hardware and associated drivers are properly installed and recognized by the Windows operating system.

- If the volume that has the deduplication storage does not have more than 12% free space available.
- If the deduplication services are not started.
- If the Windows Hotfix (Windows update) is not installed.

One of the following results is reported for each item:

Table 2-2 Environment Check results

Result	Description
Passed	There are no incompatibilities to prevent the Backup Exec installation. For hardware, this result indicates that the hardware configuration is recognized by Backup Exec.
Warning	An incompatibility with Backup Exec exists. Some of the issues may be resolved during the Backup Exec installation. A warning does not prevent Backup Exec from installing. However, if the issues are not resolved during installation, jobs may fail.
Failed	An incompatibility with Backup Exec exists, and it will cause the installation to fail. Action is required before you can successfully install Backup Exec.

Although the Environment Check runs automatically during installation, you may want to run it manually before installing Backup Exec or before backing up data with Backup Exec.

To check your environment before installing

- 1 From the installation media browser (Browser.exe), click **Preinstallation**, and then click **Backup Exec**.
- 2 Click **Next**.
- 3 Do any of the following:

To check the configuration of the local computer

Check **Local Environment Check**.

To check the configuration of a remote computer

Check **Remote Environment Check**.
- 4 Click **Next**.
- 5 If you checked **Remote Environment Check** in step 3, do one of the following, and then click **Next**:

To select the name of a computer from a list Do the following:

- Click **Add Server From List**.
- Select the computer from the list, and then click **Next**.

To add the name of a computer manually Do the following:

- Click **Add Server Manually**.
- In the **Domain** field, type the name of the domain.
- In the **Computer Name** field, type the name of the computer.
- Click **OK**.
- Type the user name and password for this computer.
- Click **OK**.

To remove the name of a computer from the list of computers on which the Environment Check runs Do the following:

- Select the computer from the list.
- Click **Remove**.

- 6 Review the results of the Environment Check, and then to save the results, click **Save Results To**.

To change the location where the Environment Check results are saved, click **Change Path** to browse to a new location.

- 7 Click **Finish**.

Microsoft SQL Server Express Edition components installed with Backup Exec

The Backup Exec installation program installs Microsoft SQL Server Express components that are required to run Backup Exec.

Backup Exec prompts you to do one of the following:

- Install the required Microsoft SQL Server Express components with Backup Exec and create a default Backup Exec instance. This is the recommended action.

Note: SQL Express can be installed on a server that runs other instances of SQL Express or full versions of SQL.

- Select a supported Microsoft SQL Server instance that already exists on the network on which you want to run Backup Exec.

If you choose to install Backup Exec into an existing supported SQL Server instance, ensure that the instance is installed before you continue with the installation.

Caution: Backup Exec may not function properly if you install it into an existing SQL Server instance that uses case-sensitive collation. It is recommended that you avoid installing Backup Exec to a SQL Server instance that uses case-sensitive collation.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

When Backup Exec applies updates, the SQL instance is stopped, which may cause other databases in the same instance to lose connectivity. If a remote SQL instance is used, ensure that Backup Exec has good network connectivity with the instance to avoid errors. A default local instance of SQL Express is recommended.

You cannot install multiple Backup Exec Databases on the same SQL Server instance.

Note: If you are installing a managed Backup Exec server, it is recommended that you select a supported local Microsoft SQL Server instance or later on which to install the Backup Exec Database for this managed server. Do not select the same SQL Server instance that is used by the central administration server.

See [“System requirements for Backup Exec”](#) on page 45.

System requirements for Backup Exec

The following are the minimum system requirements to run this version of Backup Exec:

Table 2-3 Minimum system requirements

Item	Requirements
Operating system	You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List. Note: Backup Exec supports the Backup Exec server installation on 64-bit operating systems only. You cannot install a Backup Exec server on a computer that runs the Windows Server Core installation option of Windows Server 2012 and later. You can only install the Backup Exec Agent for Windows on Server Core computers. You cannot install SQL Express or SQL Server on a Windows Server 2012 computer that is configured in a Read Only Domain Controller (RODC) role. The Read Only Domain Controller role does not let you use the local accounts that are required for SQL Express and SQL Server. When you install Backup Exec on an RODC computer you must select a remote SQL instance for the Backup Exec Database. For Windows Server 2012 and later computers, you cannot install Backup Exec or the Agent for Windows on a volume that has been enabled for data deduplication in Windows, on an ReFS volume, or on Cluster Shared Volumes.
Internet browser	Internet Explorer 9 or later, or enable Internet Explorer mode on the Microsoft Edge web browser.
Processor	Intel Pentium, Intel Core, Xeon, AMD, or compatible
Screen resolution	1024 x 768
SQL Server or SQL Express	SQL Server Express

Table 2-3 Minimum system requirements (*continued*)

Item	Requirements
Memory	Required: 8GB RAM above the operating system's requirements for the exclusive use by Backup Exec. Recommended: 16 GB RAM (or more for better performance) Note: RAM requirements may vary depending on the operations performed, the features installed, and the specific computer configuration. For the Central Admin Server feature: 8 GB RAM is required; 16 GB RAM is recommended. Recovery Disk: 4-8 GB minimum (dedicated) for the multi-lingual version. Virtual Memory Recommendations: Based on Microsoft recommendations, twice or thrice the installed RAM. For example, if the RAM is 16 GB then the virtual memory must be in the range of 32 GB to 48 GB. Refer to your Microsoft Windows documentation for instructions on how to view or set the paging file size.
User interface language	The Backup Exec user interface displays in the format that is configured in the Region and Language settings in the Control Panel. You can change the Backup Exec display language so that you can view user interface items in a different language. If Backup Exec does not support a language, the user interface displays in English. The user interface also displays in English if the menu and dialog boxes option is set to a language other than the language you want to display the user interface. If you do not use one of the supported languages listed in Table 2-4, you must install the English language pack in Windows.
Installation disk space	2 GB and more (Typical installation) 3 GB and more (Includes all features) Note: Disk space requirements may vary depending on the operations performed, the features installed, and the specific system configuration. The Backup Exec Database and catalogs require additional space. An additional 525 MB is required for SQL Express. Any disk storage that you use also requires additional space.
Other Hardware	The following hardware is recommended: ■ Network interface card or a virtual network adapter card ■ CD/DVD drive ■ A mouse

Table 2-3 Minimum system requirements (*continued*)

Item	Requirements
Storage Hardware	You can use storage media drives, robotic libraries, removable storage devices, and non-removable hard drives. You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List. For more information about support for additional drives within a robotic or virtual tape library, refer to the Backup Exec Price and Licensing Guide.

See [“Requirements for the Deduplication feature”](#) on page 972.

Table 2-4 Supported languages for the Backup Exec user interface

Language	Language code
Chinese (Simplified)	ZH
English	EN
French	FR
German	DE
Japanese	JP
Spanish	ES

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 48.

Installing Backup Exec by using the Installation Wizard

The Backup Exec installation program provides two methods of installation: typical and custom. A typical installation is simpler than a custom installation and will install all agents and features included with your license. A typical installation may be appropriate for small, simple environments. A custom installation is designed for large or complex environments, such as a remote Backup Exec server or an environment which uses the Enterprise Server feature. In a custom installation, you can choose which features and agents included with your license will be installed. This makes it ideal for a user who wants to view and control which agents and features will be installed.

In a typical installation, Backup Exec makes the following decisions for you, based on common installation scenarios:

- Backup Exec is installed to a local Backup Exec server.
- SQL Express is installed with the default instance.
- All agents and features included with your license are installed.
- Arctera Update runs automatically.

Note: Before you install, make sure that your licenses for the Backup Exec editions that you want to install are available. You must enter a license in order to install any edition of Backup Exec, but you can install the trial edition without a license.

Choose either the typical installation or the custom installation:

[How to perform a typical installation](#)

[How to perform a custom installation](#)

How to perform a typical installation

Follow these steps to perform a typical installation of Backup Exec. A typical installation of Backup Exec installs all of the features included with your license.

To install a typical installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the required version of Microsoft .NET Framework is not already installed on this computer, Backup Exec installs it.

The Backup Exec installation program uses the Microsoft .NET Framework version 4.8. However, not all versions of Windows support .NET Framework 4.8. If the Backup Exec installation program encounters an operating system that requires the use of a different version of the .NET Framework, Backup Exec blocks the installation and provides an error message that instructs you to install the required version of .NET Framework.

- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement** and click **Next**.
- 3 On the **Installation Type** panel, click **Typical installation**, and then click **Next**.
- 4 The Backup Exec Environment Check will run automatically.
- 5 Review the results of the Environment Check. Do one of the following:
 - If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.

- If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.

6 On the **Add Licenses** panel, use one of the following methods to add licenses.

To enter entitlement IDs manually

Note: If you do not have an Internet connection, import the license files manually to the Backup Exec server. To download the license file go to the Licensing Portal and then import to the Backup Exec server.

Do the following in the order listed:

- In the **Enter an Entitlement ID** field, type the appropriate entitlement ID from your license certificate.
- Click **Add to List**.
- Repeat for each entitlement ID.
- When you are finished entering the entitlement IDs, click **Next**.
- Enter the Veritas User Account credentials and then click **Download** to connect to the Licensing Portal and download the license files.

To import licenses from the License File

Do the following in the order listed:

- Click **Import License File**.
- Browse to the location of your license files, and then select the appropriate file.
- Click **Next**.

To install a trial version

Do not enter an entitlement ID or import a license file. Go to the next step.

7 If you entered entitlement IDs, do one of the following on the **Review Licenses** panel.

To install a licensed version of Backup Exec, perform the following in order:

- In the **Select a Backup Exec edition license to install on the computer** field, select the Backup Exec edition to install.
- Check the check boxes for the agents or features you want to install.
- Click the drop-down menu, and then select the number of licenses that you want to install.

To install a trial edition of Backup Exec, in the **Select a Backup Exec edition license to install on the computer** field, select **Trial**. This option is only available when you have installed a license. If you have not installed a license, a trial version will automatically be installed when you click **Next**.

When you are finished on the **Review Licenses** panel, click **Next**.

- 8 On the **Service Account** panel, provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

You cannot install Backup Exec with an account that has a blank password on a supported Windows Server unless Windows is configured to allow it. If you try to do so, the following error message appears when Backup Exec services are created: The account name and/or password supplied is not valid. Re-enter the login information and try again.

You can, however, configure Windows to allow for blank passwords. For more information, see your Windows documentation.

- 9 If you want to change the directory where Backup Exec files are installed, click **Change**, and then select a new location.

If you change the directory to a new location, ensure that you select a secure location where you can store sensitive data such as passwords.

If, during Backup Exec installation, the installer detects Backup Exec database (BEDB) files from an earlier installation, the installer provides you the option to use either the new database files or the existing database files. If you choose to use the new database files, the old database files are copied to a different location.

However, if you choose to use the existing database files, the installer warns you that the installation might fail if the existing database files are of a different version than the Backup Exec version you are trying to install.

Click **Next**.

- 10 If you want to change the directory where Backup Exec catalog is created, click **Change**, and then select a new path.

Ensure that you have the required amount of space on the folder that you select to create the catalog.

Catalog cannot be created on encrypted, compressed, or network drives.

Click **Next**.

- 11 If the **SQL Express Setup** panel appears, perform the following steps to identify the location of the SQL Express setup file:

- Click **Browse**.
- Navigate to the location where you downloaded the SQL Express 2008 R2 SP2 setup file.
- Click **OK**.

- Click **Next**.

12 On the **Remote Computers** panel, do one of the following.

To install the Agent for Windows on one remote computer, perform the following in order:

- Click **Add**.
- Select **Add a Single Computer**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.
- Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.
You must use Administrator credentials.
- Click **Next**.
- In the **Destination Folder** field, enter the path where you want to install the files.
- Click **Next**.
- After all of the computers in the list are validated and the list is complete, click **Next**.

To install the Agent for Windows on multiple computers using the same settings, perform the following in order:

- Click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.
Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push-install the Agent for Windows.
- Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.
You must use Administrator credentials.
- Click **Next**.
- In the **Destination Folder** field, enter the path where you want to install the files.
- Click **Next**.

- After all of the computers in the list are validated and the list is complete, click **Next**.

To proceed without push-installing the Agent for Windows, click **Next**.

- 13 On the **Data Backup** panel, select a location to store a copy of the existing Backup Exec database.

You can use this copy if the upgrade to Backup Exec fails.

After the upgrade to Backup Exec is complete, the deduplication storage is converted to a newer version. The deduplication storage folder remains offline till the conversion is complete. After the conversion is complete, and services are restarted, the deduplication storage comes online.

Jobs that are targeted to deduplication storage fail during the conversion process. Jobs that are scheduled on other storage continue to run during the deduplication storage conversion. The estimated time for the conversion is displayed in months, days, and hours. It is recommended that you keep a secondary copy of the deduplication data before the upgrade starts, which can be used if the conversion process fails.

Note: This panel is displayed only when you upgrade Backup Exec.

- 14 Select the **I have read and understood the information** check box, and then click **Next**.
- 15 Review the Backup Exec installation summary, and then click **Install**.
- 16 If you installed the Agent for Windows on remote computers, on the **Remote Installation** dialog box, click **Next**.
- 17 When the installation is complete, you can choose to restart the system, view the readme, or remove the Backup Exec shortcut from the desktop.
- 18 Click **Next**, and then click **Finish** to exit the wizard.

If you chose to restart the system, the computer will restart automatically.

See [“About using a self-protect backup job to protect server metadata”](#) on page 230.

The installation process creates an installation log named `BKUPINST25.htm` in the following directory on the computer where Backup Exec is installed.

For Windows Server 2012 and later: %programdata%\Veritas\Backup Exec\Logs

See [“About the installation log”](#) on page 98.

How to perform a custom installation

Follow these steps to install a custom installation of Backup Exec. A custom installation allows you to choose which agents and features will be installed, based on the licenses which you enter. Note that you cannot use a custom installation to install more agents and features than are included with your license.

To install a custom installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the required version of Microsoft .NET Framework is not already installed on this computer, Backup Exec installs it.

The Backup Exec installation wizard uses the Microsoft .NET Framework version 4.8. However, not all versions of Windows support .NET Framework 4.8. If the Backup Exec installation program encounters an operating system that requires the use of a different version of the .NET Framework, Backup Exec blocks the installation and provides an error message that instructs you to install the required version of .NET Framework.

- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement** and click **Next**.
- 3 On the **Installation Type** panel, select **Custom installation**, and then click **Next**.
- 4 On the **Menu** panel, select **Local Installation**, and then select **Install Backup Exec software and features**. Click **Next**.
- 5 The Backup Exec Environment Check will run automatically.
- 6 Review the results of the Environment Check. Do one of the following:
 - If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.
 - If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.
- 7 On the **Add Licenses** panel, use one of the following methods to enter licenses:

To enter entitlement IDs manually

Note: If you do not have an Internet connection, import the license files manually to the Backup Exec server. To download the license file go to the Licensing Portal and then import to the Backup Exec server.

Do the following in the order listed:

- In the **Enter an Entitlement ID** field, type the appropriate entitlement ID from your license certificate.
- Click **Add to List**.
- Repeat for each entitlement ID.
- When you are finished entering the entitlement IDs, click **Next**.
- Enter the Veritas User Account credentials and then click **Download** to connect to the Licensing Portal and download the license files.

To import licenses from the License File

Do the following in the order listed:

- Click **Import License File**.
- Browse to the location of your license files, and then select the appropriate file.
- Click **Next**.

To install a trial version

Do not enter an entitlement ID or import a license file. Go to the next step.

8 If you entered entitlement IDs, do one of the following on the **Review Licenses** panel:

To install a licensed version of Backup Exec, perform the following in order:

- In the **Select a Backup Exec edition license to install on the computer** field, select the Backup Exec edition to install.
- Check the check boxes for the agents or features you want to install.
- Click the drop-down menu, and then select the number of licenses that you want to install.
- If you do not make a selection, Backup Exec will select the license with the largest feature set by default.

To install a trial version of Backup Exec, select **Trial** in the **Select a Backup Exec edition license to install on the computer** field. This option is only available when you have installed a license. If you have not installed a license, a trial version will automatically be installed when you click **Next**.

The **Review Licenses** panel also allows you to view and customize the capacity for each entitlement ID.

When you are finished on the **Review Licenses** panel, click **Next**.

- 9 On the **Configure Features** panel, the **Select features to install** panel displays all of the features and agents included with the licenses for which you entered entitlement IDs. Select the check box next to any feature or agent that you want to install. You can deselect agents and features to prevent them from being installed now. Features and agents are organized by edition, based on the licenses that you have entered.

Agents and features for which you have not entered a license will be displayed, but will be unavailable. If you install a trial license, then install a license with fewer agents or features than are included with the trial, those agents and features will be removed.

When you are finished configuring agents and features, click **Next**.

- 10 If you want to install Backup Exec for any additional languages, select the language on the **Choose Languages** panel, and then click **Next**.
- 11 On the **Destination** panel, review the disk space requirements for the items that you selected to install. If you want to change the directory where the Backup Exec files are installed, click **Change**, and then select a new directory or create a new folder. It is recommended that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec will be uninstalled.

If the installer detects Backup Exec database (BEDB) files from an earlier installation, the installer provides you the option to use either the new database files or the existing database files. If you choose to use the new database files, the old database files are copied to a different location.

However, if you choose to use the existing database files, the installer warns you that the installation might fail if the existing database files are of a different version than the Backup Exec version you are trying to install.

Click **Next** when you are finished reviewing the destination information.

- 12 If you want to change the directory where Backup Exec catalog is created, click **Change**, and then select a new path.

Ensure that you have the required amount of space on the folder that you select to create the catalog.

Catalog cannot be created on encrypted, compressed, or network drives.

Click **Next**.

- 13 On the **Service Account** panel, provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

You cannot install Backup Exec with an account that has a blank password on a supported Windows Server unless Windows is configured to allow it. If you try to do so, the following error message appears when Backup Exec services are created:

The account name and/or password supplied is not valid. Re-enter the login information and try again.

You can, however, configure Windows to allow for blank passwords. For more information, see your Windows documentation.

- 14 On the **Choose SQL Server** panel, do one of the following to select a location to store the Backup Exec Database.

Note: The **Choose SQL Server** panel does not appear for upgrades. You cannot change the database location during the upgrade process. If you want to change the database location after the upgrade, use BE Utility.

To create a local Backup Exec SQL Express instance, do the following in the order listed:

- Click **Create a local Backup Exec SQL Express instance to store the Backup Exec database**.
- To change the location of the Backup Exec SQL Express instance, click **Browse**.
- Select the location, and then click **OK**.

To use an existing SQL Server 2008 R2 SP2 instance, do the following in the order listed:

- Click **Use an existing instance of SQL Server 2008 R2 with Service Pack 2 or a later SQL Server version**.
- Select the instance.
- Note that when Backup Exec is installed into an existing instance, the automated Master database restore feature is not available. To recover the Master database, replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up. For more information, See "[Microsoft SQL Server Express Edition components installed with Backup Exec](#)" on page 44.

Click **Next** when you are finished.

- 15 Backup Exec will now attempt to connect to the instance.
 - 16 If the **SQL Express Setup** panel appears, perform the following steps to identify the location of the SQL Express setup file:
 - Click **Browse**.
 - Navigate to the location where you downloaded the SQL Express 2008 R2 SP2 setup file.
 - Click **OK**.
 - Click **Next**.
 - 17 If you are prompted, select how the **Device Driver Installer** should install device drivers for the tape storage devices that are connected to the server, and then click **Next**.

It is recommended that you select **Use device drivers for all tape devices**.
 - 18 If you are prompted, enter information or choose settings for the additional features that you want to install, and then click **Next** after each selection.
 - 19 Review the Backup Exec installation summary, and then click **Install**.

The installation process takes several minutes to complete. During the process, the progress bar may not move for several minutes.
 - 20 When the installation is complete, you can choose to restart the system, view the readme, or remove the Backup Exec shortcut from the desktop.
 - 21 Click **Next**, and then click **Finish** to exit the wizard.

If you chose to restart the system, the computer will restart automatically.
- See [“About using a self-protect backup job to protect server metadata”](#) on page 230.

Installing additional agents and features to the local Backup Exec server

You can install agents and features when you install Backup Exec. However, if you have already installed Backup Exec and want to install additional agents or features, review the documentation for these features to ensure that your system meets all of its minimum requirements. In addition, you must have a valid license for any agents or features that you want to install. The Backup Exec services may be stopped while the additional features are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

Note: If the Central Admin Server feature is installed, and you want to install additional features on a managed Backup Exec server, you can pause the managed Backup Exec server. When a managed Backup Exec server is paused, the administration server does not delegate jobs to it. When the installation is complete, un-pause, or resume, the managed Backup Exec server.

See [“Pausing or resuming a managed Backup Exec server”](#) on page 1370.

To install additional Backup Exec features to the local Backup Exec server

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Features and Licenses on this Backup Exec Server**.

You may be prompted to insert the installation media.

- 2 Do one of the following:

To enter serial numbers manually

Note: If you do not have an Internet connection, import the license files manually to the Backup Exec server. To download the license file go to the Licensing Portal and then import to the Backup Exec server.

Do the following in the order listed:

- In the **Enter a Serial number** field, type the appropriate serial number from your license certificate.
- Click **Add to List**.
- Repeat for each serial number.
- Click **Next**.
- Enter the Veritas User Account credentials and then click **Download** to connect to the Licensing Portal and download the license files.

To import licenses from the license file

Do the following in the order listed:

- Click **Import License File**.
- Browse to the location of your license files, and then select the appropriate file.
- Click **Next**.

To install a trial version

Do not type a serial number or import a license file. Go to step 5.

- 3 If you entered product activation serial numbers, on the **Review Licenses** panel, the editions for which you have entered licenses will be shown. You can also edit the **Allocated Capacity** field for each edition to specify how much of your available capacity each will be able to use.

To install a trial edition of Backup Exec, select **Trial** in the **Select a Backup Exec edition license to install on the computer** field. This option is only available when you have installed a license. If you have not installed a license, a trial version will automatically be installed when you click **Next**.

- 4 Click **Next**.
- 5 Do the following:
 - Check the check boxes for the additional features that you want to install.
 - Uncheck the check boxes for the features that you want to remove.
- 6 Click **Next**.
- 7 If you are prompted, enter information or choose settings for the additional features that you want to install. Click **Next** after each selection.
- 8 Review the Backup Exec installation summary, and then click **Install**.

The Backup Exec services are stopped while the additional features are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services are restarted.

- 9 Click **Finish**.

Push-installing Backup Exec to remote computers

You cannot push install Backup Exec in the following scenarios:

- Push install from a 64-bit operating system to a 32-bit operating system
- Push install from a 32-bit operating system to a 32-bit or a 64-bit operating system

If you install Backup Exec through Terminal Services and the installation media is on a shared drive (network share) you must use a UNC path. Installation by mapped drives is not supported.

You can set up multiple server installations. Backup Exec processes up to five remote computer installations concurrently.

Note: Backup Exec installs the required version of Microsoft .NET Framework if it is not already installed on the computer where you want to push-install Backup Exec. The Backup Exec installation program uses the Microsoft .NET Framework version 4.8. However, not all versions of Windows support .NET Framework 4.8. If the Backup Exec installation program encounters an operating system that requires the use of a different version of the .NET Framework, Backup Exec blocks the installation and provides an error message that instructs you to install the required version of .NET Framework.

Before you install Backup Exec to remote computers, you should review the special considerations.

Table 2-5 Special considerations for installing Backup Exec to remote computers

Item	Consideration
Windows Server 2012 and later	To push-install Backup Exec to a computer that runs Windows Server 2012 and later, you must enable the following items on the destination computer's Windows Firewall Exceptions list ■ File and Printer Sharing ■ Windows Management Instrumentation (WMI) For more information, refer to your Microsoft Windows documentation. You cannot install Backup Exec on a volume that has been enabled for data deduplication in Windows, on an ReFS volume, or on Cluster Shared Volumes.
Symantec Endpoint Protection (SEP) 11.0 or later	To push-install Backup Exec to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. The file and printer sharing feature is turned off by default.

Note: You can also use Microsoft's Add or Remove Programs utility to install Backup Exec to a remote computer. See your Microsoft documentation for more information.

The installation process creates an installation log named `BKUPINST25.htm` in the following directory on the computer where Backup Exec is installed.

For Windows Server 2012 and later: `%programdata%\Veritas\Backup Exec\Logs`

To push-install Backup Exec to remote computers

1 Do one of the following:

To push-install Backup Exec to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Install Products**, and then click **Backup Exec**.
- On the **Welcome** panel, click **Next**.
- Select **I accept the terms of the license agreement**, and then click **Next**.
- Select **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.

To push-install Backup Exec to remote computers from the Backup Exec server

Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Agents and Backup Exec Servers on Other Servers**.

2 On the **Remote Computers** panel, do one of the following:

To install Backup Exec on one remote computer

Do the following in the order listed:

- Click **Add**.
- Select **Add a Single Computer**.
- Select **Backup Exec**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.

To install Backup Exec on multiple computers using the same settings

Do the following in the order listed:

- Click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Select **Backup Exec**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.
Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push install the features.

- 3 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 4 Click **Next**.

- 5 Select one of the following methods to enter licenses:

To enter serial numbers from your license certificate

Note: If you do not have an Internet connection, import the license files manually to the Backup Exec server. To download the license file go to the Licensing Portal and then import to the Backup Exec server.

Do the following in the order listed:

- In the **Enter a Serial number** field, type the appropriate serial number from your license certificate.
- Click **Add to List**.
- Repeat for each license for each feature or agent that you want to install.
- Click **Next**.
- Enter the Veritas User Account credentials and then click **Download** to connect to the Licensing Portal and download the license files.

To import licenses from a License File	Do the following in the order listed: ■ Click Import License File. ■ Browse to the location of your license files, and then select the appropriate file. ■ Click Next.
To install a trial version	Do not type a serial number or import a license file. Go to step 8.

6 If you entered product activation serial numbers, on the **Review Licenses** panel, do one of the following:

To install a licensed version of Backup Exec	Do the following in the order listed: ■ In the Select a Backup Exec edition license to install on the computer field, select the Backup Exec edition to install. ■ Check the check boxes for the agents or features you want to install. ■ Click the drop-down menu, and then select the number of licenses that you want to install.
To install a trial version	In the Select a Backup Exec edition license to install on the computer field, select Trial .

- 7 Click **Next**.
- 8 On the **Configure Features** panel, select any additional features that you want to install.
- For example, you can select additional standard features, or you can select the agents or the features that are available for a trial installation.

Note: When push installing Backup Exec from one server to another, the feature mapping in the Configure Feature Window will display the server from which the push install is being initiated, not the server to which the product is being push installed.

- 9 In the **Destination Folder** field, enter the location where you want to install Backup Exec.

- 10 Click **Next**.
- 11 Complete the service account credentials options as follows:

User Name	Type the user name for an Administrator account that the Backup Exec services can use. If the remote computer is in a domain, use a domain administrator's account or an equivalent account that is part of the domain administrator's group. If the remote computer is in a workgroup, use an administrator's account or an equivalent account that is part of the administrator's group on the computer.
Password	Type the password for an administrator account that the Backup Exec services can use.
Domain	If the computer is in a domain, select the domain in which the computer is located. If the computer is in a workgroup, select the computer name.

- 12 Click **Next**.
- 13 Do one of the following to select a location on which to store the Backup Exec Database, and then click **Next**.

To create a local Backup Exec SQL Express instance	Do the following in the order listed: ■ Click Create a local Backup Exec SQL Express instance to store the Backup Exec database. ■ To change the location of the database, type the new location in the Destination Folder field.

To use an existing SQL Server 2008 R2 SP2 instance

Do the following in the order listed:

- Click **Use an existing instance of SQL Server 2008 R2 with Service Pack 2 or a later SQL Server version.**
- Select the instance.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

See [“Microsoft SQL Server Express Edition components installed with Backup Exec”](#) on page 44.

Backup Exec attempts to connect to the instance.

This step is skipped during upgrades.

- 14 Click **Next**.
- 15 Make a selection for tape device drivers, and then click **Next**.

Note: You do not need to install tape device drivers if Backup Exec runs on Windows Server 2012 and later. Kernel-mode drivers and tapeinst.exe are no longer installed if Backup Exec runs on Windows Server 2012 and later.

- 16 Click **Next**.
- 17 If you are prompted, enter information or choose settings for additional features that are being installed, and then click **Next** or **OK** after each selection.
- 18 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer	Click Add , and then click Add a Single Computer .
To manually add multiple remote computers	Click Add , and then click Add Multiple Computers with the Same Settings .

To add multiple remote computers by importing an existing list of computers

Click **Import and Export**, and then select one of the following options:

- Select **Import from File** to enable Backup Exec to add the names of the remote computers from a selected list.
- Select **Import Servers Published to this Backup Exec server** to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server.

You must enter remote computer logon credentials for the list of remote computers.

To change the product that you selected to install or to change other properties you selected for this installation

Select the remote computer that you want to change, and then click **Edit**.

To delete a remote computer from the list

Select the remote computer that you want to delete, and then click **Delete**.

To save this list of remote computers and the associated remote computer logon credentials

Verify that **Save the server list for future remote install sessions** is checked.

This option enables the names and the credentials of all of the remote computers to be added automatically the next time you install Backup Exec or features to these remote computers.

To save the list of remote computers to an XML file

Click **Import and Export**, and then click **Export to File**.

You can select the location to save the Push_Export.xml file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.

To fix the errors that were located during the validation

Right-click the name of the computer, and then click **Fix Error**.

To enable Backup Exec to attempt to re-validate an invalid remote computer

Right-click the name of the computer, and then click **Retry Validation**.

- 19 After all of the computers in the list are validated and the list is complete, click **Next**.

- 20 Review the Backup Exec installation summary, and then click **Install**.

See [“About the installation log”](#) on page 98.

- 21 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Methods for installing the Agent for Windows

You can install the Agent for Windows by using the following methods, depending on your environment:

- Push-install the Agent for Windows to one or more remote computers from the Backup Exec server.
See [“Push-installing the Agent for Windows to remote computers”](#) on page 68.
- Add the remote computer to the list of servers and install the Agent for Windows on the remote computer.
See [“Adding servers that you want to back up to the list of servers on the Backup and Restore tab”](#) on page 151.
- Use a Microsoft Active Directory network to centrally manage the installation of the Agent for Windows to computers in the network.
See [“Installing the Agent for Windows in an Active Directory network”](#) on page 75.
- Use a command prompt.
See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 79.
- Use command script files.
See [“Using a command script to install the Agent for Windows”](#) on page 83.

Push-installing the Agent for Windows to remote computers

You can push-install the Agent for Windows to remote computers from a Backup Exec server. Push installations save time by eliminating the need for local access at the target computer for the installation to be successful. You can push-install the Agent for Windows to an unlimited number of remote computers. Backup Exec can process up to five active push-installations at a time.

Review the following special considerations before you install the Agent for Windows on remote computers.

Table 2-6 Special considerations for installing the Agent for Windows

Item	Consideration
ForceGuest configuration	You cannot push-install the Agent for Windows when the remote computer is in the ForceGuest configuration and it is not in a domain. ForceGuest is an operating system configuration that limits incoming users to Guest-level access. Instead, use the installation media or the network to install the Agent for Windows on the Windows computer. You can also turn off ForceGuest. Refer to your Microsoft Windows documentation for more information. See “Installing Backup Exec using the command line (silent mode)” on page 90.
Server core option of Windows Server 2012 and later	Backup Exec installs a command-line version of the Agent for Windows on the computers that run the Server Core installation option of Windows Server 2012 and later. The Backup Exec Agent Utility command-line applet is installed with the Agent for Windows. This applet lets you monitor Backup Exec operations on the remote computer. See “Backup Exec Agent Utility Command Line Applet switches” on page 964.
Windows data deduplication, ReFS volumes, and Cluster Shared Volumes	You cannot install the Agent for Windows on a volume that has been enabled for data deduplication in Windows, on an ReFS volume, or on Cluster Shared Volumes.

Table 2-6 Special considerations for installing the Agent for Windows
(continued)

Item	Consideration
Windows 8/Windows Server 2012 or later	To push-install Backup Exec features to a computer that runs Windows 8/Windows Server 2012 or later, you must enable certain items on the destination computer's Windows Firewall Exceptions list. You must enable the following items: ■ File and Printer Sharing ■ Windows Management Instrumentation (WMI) For more information, refer to your Microsoft Windows documentation. To push-install to a computer that runs the supported Backup Exec server, the destination computer must be part of a domain. For more information, refer to the Microsoft knowledge base.
Symantec Endpoint Protection 11.0 or later	To push-install features to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. File and printer sharing is turned off by default.
Trust the Backup Exec server and remote computer	When you connect to a remote computer from the Backup Exec server, you must establish a trust between the Backup Exec server and the remote computer to ensure secure communication. To establish the trust, you must add the remote computer to the list of servers on the Backup and Restore tab. See “About the list of servers on the Backup and Restore tab” on page 150.

The installation process creates an installation log named `BKUPINST25.htm` on the computer where Backup Exec is installed, and also creates an installation log named `RAWSinst25.htm` on the remote computer.

See [“About the installation log”](#) on page 98.

If there are problems installing the Agent for Windows using this method, you can try to manually install the Agent for Windows.

See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 79.

To push-install the Agent for Windows to remote computers

1 Do one of the following:

To push-install the Agent for Windows to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Installation**, and then click **Backup Exec**.
- On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- Click **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.

To push-install the Agent for Windows to remote computers from the Backup Exec server

Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Agents and Backup Exec Servers on Other Servers**.

2 Do one of the following

To install the Agent for Windows on one remote computer

Do the following steps in the order listed:

- On the **Remote Computers** panel, click **Add**.
- Select **Add a Single Computer**.
- Select **Agent for Windows**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.

To install the Agent for Windows on multiple computers using the same settings

Do the following steps in the order listed:

- On the **Remote Computers** panel, click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Select **Agent for Windows**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.

Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push install the features.

- 3 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 4 Click **Next**.
- 5 In the **Destination Folder** field, enter the path where you want to install the files.
- 6 Click **Next**.
- 7 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer

Click **Add**, and then click **Add a Single Computer**.

To manually add multiple remote computers

Click **Add**, and then click **Add Multiple Computers with the Same Settings**.

To add multiple remote computers by importing an existing list of computers

Click **Import and Export**, and then select one of the following options:

- Select **Import from File** to enable Backup Exec to add the names of the remote computers from a selected list.
- Select **Import Servers Published to this Backup Exec server** to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server.

You must enter remote computer logon credentials for the list of remote computers.

To change the product that you selected to install or to change other properties you selected for this installation

Select the remote computer that you want to change, and then click **Edit**.

To delete a remote computer from the list

Select the remote computer that you want to delete, and then click **Delete**.

To save this list of remote computers and the associated remote computer logon credentials

Verify that **Save the server list for future remote install sessions** is checked.

This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or features to these remote computers

To save the list of remote computers to an XML file

Click **Import and Export**, and then click **Export to File**.

You can select the location to save the XML file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.

To fix the errors that were located during the validation

Right-click the name of the computer, and then click **Fix Errors**.

To enable Backup Exec to attempt to re-validate an invalid remote computer

Right-click the name of the computer, and then click **Retry Validation**.

- 8 After all of the computers in the list are validated and the list is complete, click **Next**.

- 9 Review the Backup Exec installation summary, and then click **Install**.

See [“About the installation log”](#) on page 98.

- 10 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Installing updates to the Agent for Windows on remote computers

When a Backup Exec server is updated with patches, an alert is generated to warn you that the Agent for Windows on remote computers must be updated with the same patches. Additionally, in the properties for the remote computer, the property **Do the updates installed on this server match the updates installed on the backup server** indicates whether the remote computer is up to date with the Backup Exec server. From the Backup Exec console, you can update the remote computers immediately, at a scheduled time, or on a recurring schedule. You can also update a group of remote computers together.

See [“Updating Backup Exec with Arctera Update”](#) on page 101.

To install updates for the Agent for Windows

- 1 On the **Backup and Restore** tab, right-click the remote computer or the group that needs to be updated.
- 2 Select **Update**.

- 3 On the **Install Updates** dialog box, select the option for when you want to install the updates.

Recurrence	Select this option to create a recurring schedule for the job.
Recurrence Pattern	If you choose to make the job recur on a schedule, configure the frequency with which the job recurs. You can select to run the job in hourly, daily, weekly, monthly, or yearly increments.
Starting on	Enter the date on which you want the schedule to take effect.
Calendar	View all the scheduled jobs on the calendar to check for scheduling conflicts.
Keep the job scheduled for x hours before it is rescheduled	Specify the maximum amount of time after the scheduled start time at which you want Backup Exec to consider the job to be missed and reschedule it.
Cancel the job if it is still running x hours after its scheduled start time	Specify the amount of time after the job's scheduled start time at which you want Backup Exec to cancel the job if it is still running.
Run now with no recurring schedule	Select this option to run the job immediately without scheduling any more instances of it for the future.
Run on	Select a specific date on which to run the job without scheduling any more instances of it for the future.
Restart the computer automatically after installing the updates to the Backup Exec Agent for Windows when a restart is required	Select this option to enable Backup Exec to automatically restart the remote computer, if required.

- 4 Click **OK**.

Installing the Agent for Windows in an Active Directory network

You can centrally manage the installation of the Backup Exec Agent for Windows to computers in an Active Directory network. You configure the installation once, and then use a Group Policy Object to assign that installation to computers in an

Organizational Unit. The features are installed automatically whenever a computer in the Organizational Unit is started.

Note: Review your organization's deployment plans before you implement a rollout of the Backup Exec Agent for Windows to client computers. You should also review your Group Policy Desktop Management and Active Directory documentation.

Table 2-7 How to install the Agent for Windows in an Active Directory Network

Action	Description
Create a transform for the Agent for Windows. See the section called "Creating a transform" on page 77.	A transform contains the changes that you want to make to the Windows Installer package for the Agent for Windows when a computer starts, such as the installation path. Requirements to create a transform are as follows: ■ The computer on which you want to create the transform must have Microsoft Windows Server 2012 or later. ■ Any Windows Server 2012 computers on which you want to install the Agent for Windows must be running .NET Framework 2.0 SP2. ■ The computers on which you want to install the Agent for Windows must be running MSXML 6.0 SP2. ■ The computers on which you want to install the Agent for Windows must be running Microsoft Visual C++ Runtime 8.0/9.0/10.0/11.0. You can configure a Group Policy object to deploy all of the Microsoft Visual C++ Runtime components or install them manually on each Agent for Windows computer. For more information about configuring a Group Policy object, see your Microsoft Windows documentation. ■ Only assignment to computers is supported. Assignment to users is not supported.

Table 2-7 How to install the Agent for Windows in an Active Directory Network *(continued)*

Action	Description
Create a distribution point (share) that contains the source file of the Agent for Windows that you want to install. See the section called “Creating a software distribution point (share)” on page 78.	You must copy the transform that you create, and the Backup Exec RAWS32 or RAWSX64 directory, to the distribution point.
Configure a Group Policy Object to assign the transform and the RAWS32 or RAWSX64 directory in the distribution point to computers in an Active Directory Organizational Unit. See the section called “Configuring a Group Policy Object” on page 78.	The software is installed automatically when the computers in the Organizational Unit are started.

Creating a transform

To create a transform

- Do one of the following:
 - From the Backup Exec installation media browser, click **Install Products**, and then click **Agent for Windows**.
 - From a Backup Exec server on which Backup Exec is installed, go to <Backup Exec install path>\Backup Exec\Agents\RAWS32 or RAWSX64 and double-click **Setup.exe**.
- On the **Welcome** panel, click **Next**.
- On the **Install Type** panel, click **Create a Transform to use Active Directory to install the Agent for Windows**, and then click **Next**.
- On the **Install Option** panel, in the **Destination Folder** area, enter the path where you want to install the files.
- Click **Next**.
- Enter a file name and a path where the transform will be created, and then click **Next**.

Use a meaningful file name for the transform. For example, the name can include the names of the features in the transform and the platform you plan to apply the transform to, such as AgentDefaultPathNoPublishing.

- 7 To create the transform, click **Install**.
- 8 After the transform is created, set up a distribution point for the source files.

Creating a software distribution point (share)

To install the Agent for Windows in an Active Directory network, you must create a software distribution point after you create a transform.

Table 2-8 How to create a software distribution point (share)

Step	Description
Step 1	Create a shared folder, and then set permissions so that the client computers that will run the installation have access to the shared folder.
Step 2	Copy the following directories from the Backup Exec server to the shared folder: ■ RAWS32 or RAW SX64 ■ MSXML ■ VCRedist ■ DotNetFx By default, these folders are located in <Backup Exec install path>\Backup Exec\Agents.
Step 3	Copy the transform from the path where it was created to the RAWS32 or RAW SX64 directory on the shared folder.
Step 4	Configure a Group Policy Object to deploy the source files.

Configuring a Group Policy Object

To install the Agent for Windows in an Active Directory network, you must configure a Group Policy Object after you create a software distribution point and create a transform.

To configure a Group Policy Object to deploy the software

- 1 From the Active Directory snap-in that manages users and groups, click **Properties**, and create a new Group Policy Object or edit an existing one.

Refer to your Microsoft Windows documentation for information on creating a Group Policy Object.
- 2 Under **Computer Configuration**, expand **Software Settings**.
- 3 Right-click **Software Installation**, click **New**, and then click **Package**.
- 4 On the **File Open** dialog box, browse to the software distribution point by using the Universal Naming Convention (UNC) name, for example, \\server name\share name, select the package file, and then click **Open**.
- 5 Select the package file **Backup Exec Agent for Windows.msi**, and then click **Open**.
- 6 When you are prompted, apply the **Advanced Option**.
- 7 After Active Directory checks the MSI package, on the **General Properties** tab, make sure that the correct versions of the features are being installed.
- 8 On the **Deployment** tab, set up the configuration for your environment.

Make sure the option **Make this 32-bit x86 application available to WIN64 machines** is not selected.

If you want the Agent for Windows to be uninstalled if the computer is removed from the Organization Unit, select the option **Uninstall this application when it falls out of the scope of management**.
- 9 On the **Modifications** tab, click **Add**, browse to the share, and select the transform that you created.
- 10 Select **Open**, and make any other changes that are necessary, and then click **OK**.
- 11 Close all of the dialog boxes.

When a computer in the Organizational Unit that you specified is started, the transform is processed and the features that you specified are installed.
- 12 View the installation log that is created on the destination computers to verify the installation of the Agent for Windows.

Using a command prompt to install the Agent for Windows on a remote computer

You can install the Agent for Windows by using a command prompt.

The installation process creates an installation log named RAWSin25.htm.

See [“About the installation log”](#) on page 98.

To use a command prompt to install the Agent for Windows on a remote computer

- 1 At a remote computer, map a drive letter to the Agents directory. By default, the Agents directory is located at the following path:

<Backup Exec install path>\Backup Exec\Agents

or you can copy the following folders to the same local directory:

To install to a 32-bit computer: RAW32, MSXML, VCRedist, and
DotNetFx folders

To install to a 64-bit computer: RAWX64, MSXML, VCRedist, and
DotNetFx folders

- 2 Open a command prompt and type the drive letter that you mapped in step 1 and the following path:

To install to a 32-bit computer: \RAW32

To install to a 64-bit computer: \RAWX64

3 Do one of the following:

To install the Agent for Windows to a 32-bit computer without publishing enabled

Run the following command:

```
setup.exe /RANT32: /S: /DISADVRT:
```

To install the Agent for Windows to a 32-bit computer with publishing enabled

Run the following command:

```
setup.exe /RANT32: /S: /ADVRT:
Backup Exec server name 1 Backup
Exec server name 2
```

To install the Agent for Windows to a 32-bit computer and restart the computer automatically

Run the following command:

```
setup.exe /RANT32: /S: /BOOT:
```

To install the Agent for Windows to a 64-bit computer without publishing enabled

Run the following command:

```
setup.exe /RAWSX64: /S: /DISADVRT:
```

To install the Agent for Windows to a 64-bit computer with publishing enabled

Run the following command:

```
setup.exe /RAWSX64: /S: /ADVRT:
Backup Exec server name 1 Backup
Exec server name 2
```

To install the Agent for Windows to a 64-bit computer and restart the computer automatically

Run the following command:

```
setup.exe /RAWSX64 /S: /BOOT:
```

The Agent for Windows is installed on the remote computer in the following directory:

If you installed the Agent for Windows to a 32-bit computer:

```
<Backup Exec install
path>\Backup Exec\RAWS32
```

If you installed the Agent for Windows to a 64-bit computer:

```
<Backup Exec install
path>\Backup Exec\RAWSx64
```

Using a command prompt to uninstall the Agent for Windows from a remote computer

You can uninstall the Agent for Windows by using a command prompt.

To use a command prompt to uninstall the Agent for Windows from a remote computer

- 1 At the remote computer, map a drive letter to the Agent for Windows directory using the following path:

To uninstall the Agent for Windows from a 32-bit computer: `<Backup Exec install path>\Backup Exec\Agents\RAWS32`

To uninstall the Agent for Windows from a 64-bit computer: `<Backup Exec install path>\Backup Exec\Agents\RAWSX64`

- 2 Open a command prompt, and then type the drive letter that you mapped in step 1.
- 3 Run the following command:

To uninstall the Agent for Windows from a 32-bit computer: `setup.exe /RANT32: /S: /U:`
The `/S:` parameter is used to run the operation in silent mode, without the benefit of a user interface. The `/U:` parameter specifies an uninstall operation.

To uninstall the Agent for Windows from a 64-bit computer: `setup.exe /RAWSX64: /S: /U:`

See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 79.

Using a command script to install the Agent for Windows

You can use command script files to install the Agent for Windows. The command script files are included in the Backup Exec installation directory.

The installation process creates an installation log named `RAWSinst25.htm`.

See [“About the installation log”](#) on page 98.

To use a command script to install the Agent for Windows

- 1 Map a drive letter to the Agents directory on a Backup Exec server. By default, the Agents directory is located at the following path:

<Backup Exec install path>\Backup Exec\Agents

- 2 Do one of the following:

To install the Agent for Windows on a 32-bit computer

In the RAW32 directory, double-click **setupaa**.

To install the Agent for Windows on a 64-bit computer

In the RAW64 directory, double-click **setupaax64**.

Using a command script to uninstall the Agent for Windows

A command script file is available to uninstall the Agent for Windows.

To use a command script to uninstall the Agent for Windows

- 1 Map a drive letter to the Agents directory on a Backup Exec server. By default, the Agents directory is located at the following path

<Backup Exec install path>\Backup Exec\Agents

- 2 Do one of the following:

To uninstall the Agent for Windows from a 32-bit computer

In the RAW32 directory, double-click **Uninstallaafo**.

To uninstall the Agent for Windows from a 64-bit computer

In the RAW64 directory, double-click **Uninstallaax64**.

- 3 Restart the remote computer.

See ["Using a command script to install the Agent for Windows"](#) on page 83.

Installing the Remote Administrator

The Remote Administrator lets you administer the Backup Exec server from a remote Windows server or workstation. To support the Remote Administrator, the Backup Exec system services must be running on the Backup Exec server that you want to administer.

Note: Backup Exec does not support the Remote Administrator on 32-bit operating systems.

You cannot use the Remote Administrator to administer the servers that have different versions of Backup Exec installed. However, you can use side-by-side installations of the Remote Administrator to manage different versions of Backup Exec.

To install the Remote Administrator

- 1 From the installation media browser, click **Install Products**.
- 2 Click **Backup Exec**.
- 3 On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- 4 On the **Installation Type** panel, select **Custom installation**, and then click **Next**.
- 5 Check **Local Installation**, and then click **Install Remote Administration Console only**.
- 6 Select the **Use the logged on user's credentials as default authentication (also applicable for smart card users)** check box if you want to use the logged on Windows user's credentials as default authentication for the Remote Administration Console to log on to the Backup Exec server.

By default, this check box is not selected and you must enter the credentials to log on to the Backup Exec Server.

This authentication can also be used if you log on to Windows using different forms of authentication such as smart cards.
- 7 Click **Next**.
- 8 On the **Destination** panel, do the following:
 - Review the disk space requirements for the installation.
 - To change the location where the files are installed, click **Change** to select another directory for the installation.
- 9 Click **Next**.
- 10 Review the Backup Exec installation summary, and then click **Install**.
- 11 Click **Next**, and then click **Finish** to exit the wizard.

See ["Running the Remote Administrator"](#) on page 88.

Installing the Remote Administrator using the command line

You can use silent mode installation to install the Remote Administrator. Options for the Remote Administrator are specified with the use of additional command switches.

Note: Backup Exec does not support the Remote Administrator on 32-bit operating systems.

You cannot use the Remote Administrator to administer the servers that have different versions of Backup Exec installed. However, you can use side-by-side installations of the Remote Administrator to manage different versions of Backup Exec.

To install the Remote Administrator using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change to the following directory:

```
\be\winnt\install\bex64
```

- 4 Type `setup /RA:` and the appropriate switches. For example:

```
setup /RA: /S:
```

The command line switches used for silent mode installation of the Remote Administrator are described in the following table.

Remember the following general rules for using these switches:

- Substitute values appropriate for your environment for values in *italics*; for example, substitute your password for *password*.
- Enclose the value in quotation marks if it contains spaces, such as "*<Backup Exec install path>\Backup Exec*".

Table 2-9 Command line switches for Remote Administrator silent mode installation

Switch	Additional Switches	Description
/RA:		Installs Remote Administrator using the options that are specified with the additional switches.

Table 2-9 Command line switches for Remote Administrator silent mode installation (continued)

Switch	Additional Switches	Description
	/DEST:"drive:\path"	Specifies the local path where Remote Administrator will be installed. Otherwise, the default path <code><Backup Exec install path>\Backup Exec</code> is used.
	/DOCS:	Installs the online documentation.
	/NOINSTALL:	Lets you select all install features without actually installing the Backup Exec software. This option can be used with the /CPF: switch.
	/CPF:"path\filename.cpf"	Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters such as the password.
-?		Provides help on all command-line operations, usage, and special switches.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 90.

Running the Remote Administrator

The Remote Administrator lets you administer the Backup Exec server from a remote Windows server or workstation. To support the Remote Administrator, the Backup Exec server requires that the Backup Exec system services must be running.

You may be prompted for a user name and password to browse some network shares even if you are logged into the Remote Administrator computer under an account that is valid for those shares. Provide a domain-qualified user name and password when prompted (for example, domain1\howard).

For workgroup accounts, when logging in between different workgroups, you can provide only a user ID when prompted, and leave the workgroup line blank.

Note: Backup Exec does not support the Remote Administrator on 32-bit operating systems. You cannot use the Remote Administrator to administer the servers that have different versions of Backup Exec installed. However, you can use side-by-side installations of the Remote Administrator to manage different versions of Backup Exec.

See [“Installing the Remote Administrator ”](#) on page 84.

To run the Remote Administrator

- 1 Click **Start**.
- 2 Point to Programs, and then click **Backup Exec**.

If you are connecting to a Remote Administration Console from a Backup Exec server, click the Backup Exec button, and then select **Connect to Backup Exec Server**.

3 Select the appropriate options.

Manage services

Select this option to access the Backup Exec Services Manager to stop and start services or to set the logon credentials that are used to run the services.

Server name

Enter the name of the Backup Exec server. You can select the name from the list or type the name of the server if you are running the Remote Administrator from a Backup Exec server.

Each server in the domain that has Backup Exec installed automatically appears in the list box.

User name

Enter an administrator user name for the server to which you want to connect. Enter the user name using the Domain\Username format.

You cannot log on to the Remote Administration Console with a user name that has a blank password. You must configure Windows to allow blank passwords. Otherwise, the error message "Logon failure: user account restriction" appears. For more information, see your Windows documentation.

Password

Enter the password for the user.

Domain

Enter the domain to which the user belongs. Enter the name of the domain that was used in the **User name** option.

Connect with the logged on user's credentials (also applicable for smart card users)

Select this check box to connect to the Remote Administration Console using the logged on Windows user's credentials.

This authentication can also be used if you log on to Windows using different forms of authentication, such as smart cards.

The default behavior of this check box depends on whether you selected this option during the Remote Administrator installation.

The status of the local services appears at the bottom of this dialog box. If you try to connect to a server and the connection fails, this dialog box displays the services status for the server you attempted to connect to.

4 Click **OK**.

If you have used the **Lock Console** feature to lock the Backup Exec console you need to enter credentials to connect to the Backup Exec server.

See [“Locking and unlocking the Backup Exec Console”](#) on page 116.

Installing Backup Exec using the command line (silent mode)

Installing Backup Exec using the command line is referred to as silent mode installation. This method of installation uses the setup.exe program on the Backup Exec installation media, a series of command switches, and the /S: switch.

Requirements for Command Line Installation include the following:

- Backup Exec installation media.
- Administrator privileges on the computer where you want to install, configure, or uninstall Backup Exec.

The installation process creates an installation log named `BKUPINST25.htm` on the computer where Backup Exec is installed.

See [“About the installation log”](#) on page 98.

To install Backup Exec using the command line (silent mode)

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change to the following directory:

```
\be\winnt\install\bex64
```

4 Type `setup /TS:` and the appropriate switches. For example:

```
setup /TS: /USER:<user> /DOM:domain /PASS:password
/SLF:C:\path\slf.slf,C:\path\slf2.slf /S:
```

See [“Command line switches for silent mode installation of Backup Exec”](#) on page 91.

If you use the command line switches without the `/S:` switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, if `/S:` had been left in the above example, the Backup Exec installation program launches with the user name, domain, password, and license appearing on the installation dialog boxes.

5 Press **Enter**.

Command line switches for silent mode installation of Backup Exec

The command line switches used for silent mode installation of Backup Exec are described in the following table.

The following are general rules for using these switches:

- Substitute values appropriate for your environment for the values that are shown in *italics*. For example, substitute a user name for *user*.
- Enclose the value in quotation marks if it contains spaces, such as "Operations Weekly Backup".

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 90.

Table 2-10 Command line switches for silent mode installation of Backup Exec

Switches	Description
<code>/S</code>	Performs a silent install.

Table 2-10 Command line switches for silent mode installation of Backup Exec (continued)

Switches	Description
/USER:user /DOM:dm /PASS:pw	Required. Specifies an existing user, domain, and password for the Backup Exec system service account. Silent mode installation does not create a user. Note: When using /PASS:, if a quote is needed as part of the password, specify it as \". For example, if the password is pass\"word, type it as /PASS:pass\"word. If the characters \" are used as part of the password, you must precede each character with a \. For example, if the password is pass\"word, type it as /PASS:pass\\\"word.
/DEST:drive:\path	Specifies the local path where Backup Exec is installed. Otherwise, the default path <i><Backup Exec install path>\Backup Exec</i> is used.
/DESTCATALOG:	Installs the Backup Exec catalog on the specified location.
/DOCS:	Installs the online documentation.
/BELANG:languagecode	Installs the Backup Exec language resource files. Specify one or more of the following switches with the Backup Exec language switch to indicate which language files you want to install: ■ EN installs English. ■ DE installs German. ■ ES installs Spanish. ■ FR installs French. ■ JP installs Japanese. ■ ZH installs Simplified Chinese. The following example shows how the /BELANG switch can be used to install English, German, and Spanish: /BELANG:"EN DE ES"
/NOINSTALL:	Lets you select all install options without installing the Backup Exec software. This option can be used with the /CPF: switch.

Table 2-10

Command line switches for silent mode installation of Backup Exec (continued)

Switches	Description
/CPF:path\filename.cpf	Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters.
/SLF:slf file location	Specifies one or more licenses to use for installing Backup Exec and additional features. Licenses are not required to install the Remote Administrator. You may specify up to 99 licenses. If none are specified, then a trial copy of Backup Exec is installed. The following examples show how the /SLF switch can be used: /SLF:C:\path\slf1.slf /SLF:C:\path\slf1.slf,C:\path\ slf2.slf, C:\path\ slf3.slf Note: If you install a license for an feature or agent, you must also type a switch that specifies the feature or agent. The switches that specify an feature or agent are included in this table.
/DISABLETELEMETRY:	Disables the option to send Backup Exec usage data over the web.
/BOOT:	Automatically initiates a restart of the computer during a silent install or uninstall.
/RA:	Installs Backup Exec Remote Administrator using the options that are specified with the additional switches.
/TD:NEW, ALL, or NONE	Note: None of the additional /TD switches are supported for Windows Server 2012 or later. /TD:NEW installs tape drivers only for the drives that do not have drivers loaded. /TD:ALL installs tape drivers for all drives. /TD:NONE does not install tape device drivers.
/DBSERVER:server\instance	Installs the Backup Exec Database to the specified SQL server.

Table 2-10 Command line switches for silent mode installation of Backup Exec *(continued)*

Switches	Description
/BACKUPDATA:	Specifies if a copy of your current Backup Exec Database should be stored during an upgrade of Backup Exec.
/BACKUPDIR:	Specifies the location to store the Backup Exec Database during an upgrade of Backup Exec. The folder you select must be empty. A copy of your current Backup Exec Database is placed in the location that you specify and is used if the upgrade fails. Note: This switch is ignored if the corresponding /BACKUPDATA: switch is not specified.
/DBINSTPATH: <i>SQL Express destination folder</i>	Installs the default instance of SQL Express in the specified folder.
/UNINSTALL:	Invokes the uninstallation process.
/ADVRT: <i>Backup Exec server name</i>	Installs the Agent for Windows and enables publishing.
/DISADVRT:	Installs the Agent for Windows without publishing it.
/PARAMS: <i>parameter file</i>	Use values in the given parameter file for installation settings.
/SQLXSETUP: <i>SQL Express Install Package</i>	Specifies the location of the language-specific install package for Microsoft SQL Server Express.
/LOADER:	Installs the Library Expansion feature. This feature is not supported for Windows Server 2012 or later.
/NFR:	Installs the Not For Resale Edition. You must enter a Not For Resale license to install this edition.
/TRIAL:	Installs the Trial Edition.
/APPLICATIONS:	Installs the Agent for Databases and Applications.
/VRAY	Installs the V-Ray Edition. You must enter a V-Ray license to install this edition.
/CAPACITY	Installs the Capacity Edition. You must enter a Capacity license to install this edition.

Table 2-10 Command line switches for silent mode installation of Backup Exec (*continued*)

Switches	Description
/CAPACITYLITE	Installs the Capacity Edition Lite. You must enter a Capacity Edition Lite license to install this edition
/VIRT:	Installs the Agent for VMware and Hyper-V.
/ENTSERVER:	Installs the Enterprise Server feature. You must use one or both of the following switches with the Enterprise Server feature switch to indicate which features you want to install. ■ /CASO: Installs the Central Admin Server feature. ■ /ADBO: Installs the Advanced Disk-based Backup feature.
/ADBO:	Installs the Advanced Disk-based Backup feature. You must use /ENTSERVER: with this switch.
/CASO:	Installs the Central Admin Server feature. You must use /ENTSERVER: with this switch.
/MMS:CAS <i>server name</i>	Creates a managed Backup Exec server for use with the Central Admin Server feature.
/CASOPVLOCAL: <1 or 0>	/CASOPVLOCAL:1 indicates that device and media data is stored locally on the managed server. Use this switch with /MMS:. /CASOPVLOCAL:0 indicates that device and media data is stored on the administration server. Use this switch with /MMS:.
/ACCESSCATALOGSANDRESTORE:	Enables unrestricted access to catalogs and backup sets for restore. This switch is used with the /MMS:<CAS server name> switch and it replaces the /SSO:<primary server name> switch.
/NTA:	Installs the Agent for Windows.
/NDMP:	Installs the NDMP feature.

Table 2-10 Command line switches for silent mode installation of Backup Exec *(continued)*

Switches	Description
/RALS:	Installs the Agent for Linux and Unix.
/DEDUPE:	Installs the Deduplication feature.
/VTL:	Installs the Virtual Tape Library Unlimited Drive feature.
/RMAL:	Installs the Remote Media Agent for Linux.
/COPYCONFIG:	Installs the Copy Server Configuration feature.
/BRONZE_CAPACITY:	Installs the Backup Exec Bronze Edition - Capacity. You must enter a Bronze Edition - Capacity license to install this edition.
/SILVER_CAPACITY:	Installs the Backup Exec Silver Edition - Capacity. You must enter a Silver Edition - Capacity license to install this edition.
/GOLD_CAPACITY:	Installs the Backup Exec Gold Edition - Capacity. You must enter a Gold Edition - Capacity license to install this edition.
/BRONZE_INSTANCE:	Installs the Backup Exec Bronze Edition - Instance. You must enter a Bronze Edition - Instance license to install this edition.
/SILVER_INSTANCE:	Installs the Backup Exec Silver Edition - Instance. You must enter a Silver Edition - Instance license to install this edition.
/GOLD_INSTANCE:	Installs the Backup Exec Gold Edition - Instance. You must enter a Gold Edition - Instance license to install this edition.
/SCFBACKUP	Installs the Microsoft 365 feature.
/SIMPLE_CORE_PACK:	Installs the Backup Exec Simple Core Pack Edition - Instance. You must enter a Simple Core Pack Edition - Instance to install this edition.
/SIMPLE_CAPACITY_CORE_PACK:	Installs the Backup Exec Simple Capacity Core Pack. You must enter a Backup Exec Simple Capacity Core Pack license to install this edition.

Table 2-10 Command line switches for silent mode installation of Backup Exec *(continued)*

Switches	Description
/RESTORE_ONLY_ED:	Installs Backup Exec Restore Only. You must enter a Backup Exec Restore Only license to install this edition.
-?	Provides help on all command-line operations, usage, and special switches.

Creating and using installation parameter files

If you use the command line switches without the /S: switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, suppose you type:

```
SETUP /TS: /USER:user /DOM:domain /PASS:password /SLF:"C:\path
name\slf1.slf"
```

The Backup Exec installation program is launched. The screens that let you enter the logon credentials and the license will appear with the information you provided on the command line.

You can also use the /CPF: command to create a parameter file that contains all of the command line options you provided. This parameter file can then be used to provide the options for installing either Backup Exec or the Remote Administrator. Note that the file is not encrypted, which exposes parameters such as the password.

To create installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change to the following directory:

```
BE\WINNT\INSTALL\Bex64
```

- 4 Type `setup /TS:` and the appropriate switches, including /CPF: and the full path name of the parameter file. For example, type:

```
setup /TS: /USER:user /DOM:domain /PASS:password /SLF:"C:\path
name\slf1.slf" /CPF:"A:\file name" /S:
```

Backup Exec will be installed on your server and a parameter file containing the user name, domain, password, and license will be saved to a removable device. You can use this parameter file to install to another computer.

To use installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change to the following directory:

`BE\WINNT\INSTALL\Be64`
- 4 Type:`SETUP /PARAMS:"A:\file name" /S:`
- 5 If you want to overwrite a parameter, specify the new parameter. For example, to change the password, type:`SETUP /PARAMS:"A:\file name" /PASS:new password /S:`

About the installation log

Backup Exec creates an installation log file, named `BKUPINST25.htm`, when you install Backup Exec and when you install patches. This log file can help you troubleshoot installation problems. The log file provides links to tech notes for the most common errors. If you install the Agent for Windows, a log file called `RAWSinst25.htm` is also created.

In addition, the text in the log file uses the following colors so you can identify warnings and errors:

Table 2-11 Installation log colors

This color	Indicates
Black	Normal operations
Orange	Warning messages
Red	Error messages

The `BKUPINST25.htm` file is located in the following location:

For Windows Server 2012 and later: `%ProgramData%\Veritas\Backup Exec\Logs`

Note: The ProgramData folder is a hidden folder. If you do not see the ProgramData folder, refer to the Microsoft Windows documentation for instructions on how to display hidden folders.

Viewing the Installation Summary Report

Backup Exec creates an Installation Summary Report that includes the configuration settings that you selected during the installation process. The Installation Summary Report is updated with the product name and entitlement IDs when you install additional agents or features. It is also updated when you remove agents or features from Backup Exec.

The Installation Summary Report is stored in the following location:

For Windows Server 2012 and later: %programdata%\Veritas\Backup Exec\Logs\InstallSummary\<computer name>.htm

The Installation Summary Report is available for viewing from the Backup Exec Administration Console or the **Home** tab at any time after the installation has completed.

To view the Installation Summary Report

◆ Do one of the following:

To view the Installation Summary Report from the Administration Console

Click the Backup Exec button, select **Installation and Licensing**, and then select **Installation Summary Report**.

To view the Installation Summary Report from the **Home** tab

On the **Home** tab, in the **Installation and Upgrades** group, click **Installation Summary Report**.

Repairing Backup Exec

If you have missing or corrupted Backup Exec files or registry keys on the local Backup Exec server, run the Repair feature. The program stops all Backup Exec services, reinstalls corrupted files and registry keys, reinstalls tape devices (standalone drives and libraries), and restarts the services. The database is not reinstalled.

Any changes that are made to Backup Exec program files and registry keys are reset to the original settings.

To repair Backup Exec

- 1 Close the Backup Exec application.
- 2 From the Windows Control Panel, select the option to uninstall a program.
- 3 Select **Backup Exec**, and then click **Change**.

- 4 Select **Local installation** and **Repair**, and then click **Next**.
Ensure that the **Remote installation** option is not selected.
- 5 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 6 Select **Install**.
You may be prompted to insert the installation media.
- 7 Click **Finish**.

About updates to Backup Exec using Arctera Update

Arctera Update provides updates to Backup Exec, and is installed automatically with Backup Exec.

Arctera Update can be run manually, or can be configured to run automatically every day at a specific time. It can also be configured to download updates automatically and raise an alert after the download is complete, or to detect updates and raise an alert instead of proceeding to download them. Arctera Update is integrated with Backup Exec, and can only be accessed from within the Backup Exec interface. If you enable the automatic update feature, you can configure Arctera Update to poll the main Arctera web server on a scheduled interval. If Arctera Update installs any files, the `BKUPINST25.htm` installation log file is updated with information about those files.

If you run Arctera Update through the Remote Administration Console (RAC), please note the following behaviors:

- Arctera Update only downloads and installs updates for the local server on which RAC is installed, not the remote media server to which the RAC is pointed.
- Any Arctera Update alerts apply to the remote media server.
- If you navigate to “Installed Updates” through the RAC interface, you will only see the updates which have been installed on the remote media server, not the updates which have been installed on the local RAC server.
- To view updates which have been installed on the local RAC server, navigate to the Windows “Programs and Features” control panel.
- If you change the Arctera Update settings while using the Remote Administration Console, those settings will be changed on the remote server.

For information about the best practices to use Arctera Update, refer to the *Backup Exec Best Practices*.

Note: During the installation and update process, the Backup Exec services are stopped and started one time during a Arctera Update session, regardless of the number of updates that are being installed. All selected patches are installed in order.

For more information about how to update Backup with Arctera update, refer to the following topic:

See [“Updating Backup Exec with Arctera Update”](#) on page 101.

Updating Backup Exec with Arctera Update

Scheduling automatic Backup Exec updates using Arctera Update

You can schedule Arctera Update to check for updates automatically at a specific time every day. By default, Arctera Update will check for updates at 10:00 PM.

At the scheduled time, Arctera Update automatically connects to the appropriate website, and then determines if your files need to be updated. You can configure Arctera Update to either download updates automatically and raise an alert when the download is complete, or to detect available downloads, raise an alert notifying you, and wait for your confirmation to download them.

Backup Exec sends the following Arctera Update alerts:

Table 2-12 Arctera Update alerts

Backup Exec sends this alert	When
Arctera Update information for Backup Exec	There are <n> updates available. To download and install the available updates, click Installation and Licensing > Arctera Update .
Arctera Update information for Backup Exec	There are <n> updates available and ready for installation. To install available updates, click Installation and Licensing > Arctera Update .
Arctera Update Error for Backup Exec	The query for available updates failed: error code=<n>

To schedule automatic update downloads using Arctera Update

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Arctera Update**.
- 3 Select **Check for updates daily**.
- 4 Enter a new time in the **At:** field to set the time of day at which Arctera Update will run. By default, this field is set to 10:00 PM.
- 5 Select one of the following options to determine what Arctera Update should do when it detects available updates:
 - If you select **Send an alert notification when updates are available, but do not download or install the update**, Arctera Update will send you an alert when it detects available updates, but will not download or install them.
 - If you select **Download updates first, and then send an alert notification**, Arctera Update will download any available updates, then send you an alert notifying you that updates are ready to be installed.
- 6 Click **OK** to close the settings window, or **Click here to run Arctera Update** to launch Arctera Update.

Running Arctera Update manually

When you launch Arctera Update, it will automatically search for and display any available updates.

To download and install updates using Arctera Update, do the following:

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Arctera Update**. Arctera Update will check for available updates and display them.
- 2 Select the check boxes next to the updates that you want to install.
- 3 Click **Install**.
- 4 In the **Start Patch Process** dialog box, click **Yes** to continue with the installation process, or **No** to cancel.
- 5 If you clicked **Yes** in the **Start Patch Process** dialog box, Arctera Update will download the update and launch the update installer. Follow the installation instructions included with the update to complete the installation process.

See [“About updates to Backup Exec using Arctera Update”](#) on page 100.

See [“Viewing installed updates”](#) on page 103.

See [“Installing updates to the Agent for Windows on remote computers”](#) on page 74.

Viewing installed updates

You can view the updates that are installed on a Backup Exec server. To do this, you must be logged on with administrator privileges. Click the Backup Exec button, select **Installation and Licensing**, and then select **Installed Updates**.

If a feature pack is installed before another feature pack, the earlier feature pack is no longer displayed because it is included with the later feature pack.

A hotfix that is offered after the feature pack is released is displayed with the previous feature pack.

Table 2-13 Installed Updates options

Item	Description
Installed Updates	Lists the hot fixes and the service packs that are installed on the Backup Exec server.
Click here to run Arctera Update	Lets you run Arctera Update to install hot fixes and service packs.
Click here to view available updates	Lets you view the hot fixes and service packs that are available for download.

Uninstalling Backup Exec updates

To uninstall hotfixes or feature packs which you have already installed, you must use the Windows **Programs and Features** control panel.

Follow the steps below to uninstall hotfixes or feature packs:

- 1 Open the Windows Start Menu.
- 2 Type "**Programs and Features**" to search for the Programs and Features control panel. Select the control panel when it appears in the search results.
- 3 In the left pane of the **Programs and Features** control panel, click **View installed updates**.
- 4 Scroll to the **Backup Exec (TM)** entry. Any hotfixes and feature packs which you have already installed will be listed here.
- 5 Select a hotfix or feature pack and click **Uninstall** to begin the uninstallation process.

Viewing license information

You can view information about the agents and features that are licensed and installed on a Backup Exec Server.

For details on updated licensing, see the [Backup Exec Licensing Guide](#).

To view license information

- ◆ Click the Backup Exec button, select **Installation and Licensing**, and then select **License Information**.

Backup Exec license contract information

After you purchase or import license contracts for Backup Exec, the Licensing Portal will be updated with your license contract information.

The entitlement associated with the ID contains the following information:

- Licensing information for the product purchased.
- Activation information.
- License information, if it was purchased.

Backup Exec uses the contract expiration information to automatically set Backup Exec alerts that remind you to renew the license contracts before they expire. Reminder alerts are set at 30-day, 60-day, and 90-day intervals, based on the expiration date of the license contract. If you do not renew the license contract, an alert is sent when the license contract expires.

Based on the expiration date of the license, Backup Exec checks for a license update in the Licensing Portal. If the license expiry period is greater than 60 days, Backup Exec checks for a license update every 30 days. If the license expiry period is less than 60 days, Backup Exec checks for a license update every 7 days. If you want to run a check for a license update in addition to this check from Backup Exec, go to the **Home** tab > **Support** group > and click **Synchronize Entitlements with Licensing Portal**.

The Licensing Information panel on the Backup Exec console displays the number of days until the currently installed license contract expires.

To view Backup Exec license contract information

- ◆ Click the Backup Exec button, select **Installation and Licensing**, and then select **License Contract Information**.

See [“Updating expired license contracts”](#) on page 105.

Updating expired license contracts

When your license contracts expire, follow these steps to update them.

Note: You cannot use the Remote Administrator on a remote Windows server or workstation to update license contracts.

Table 2-14 How to update expired license contracts

Step	Additional information
Purchase new license contracts.	Contact the reseller.
Get an updated license file.	Update your current entitlements with a new date and a new service contract number.
Launch the installation wizard from the Backup Exec Administration Console.	Use the option Install Features and Licenses on this Backup Exec Server on the Installation and Licensing menu, which is accessed from the Backup Exec button.
Use the installation wizard to add the new entitlement IDs, and then remove the expired entitlement IDs.	After you select the expired entitlement IDs from the list, use the Remove option.

See [“Backup Exec license contract information”](#) on page 104.

Managing license contract customer numbers

Backup Exec provides a place where you can store all of your license contract customer numbers.

You need to provide these numbers when you call technical support.

To manage license contract customer numbers

- Click the Backup Exec button, select **Installation and Licensing**, and then select **License Contract Customer Numbers**.
- Do one of the following:
 - To add a new customer number, click **New**, and then enter your customer number and any notes for this number.
You can find your customer number in the Entitlement Owner box on the certificate.

- To remove a customer number, select the number from the list, and then click **Delete**.

3 Click **Close**.

About upgrading to Backup Exec

In order to upgrade from older versions of Backup Exec to Backup Exec Update 1 (20.1) or later, you must provide a valid license at the time of the upgrade. Existing licenses cannot be used to upgrade to the current version, even with a valid maintenance agreement. You do not need to uninstall the previous version of Backup Exec before upgrading. The current version will be installed over the previous version. Different versions of Backup Exec cannot coexist on the same computer.

Note: Backup Exec supports the Backup Exec server installation on 64-bit operating systems only. However, you can install the Agent for Windows on 32-bit operating systems.

Most settings, all catalogs, and all data directories from previous versions of Backup Exec are kept, unless you choose to remove them. This version of Backup Exec can read and restore data from any previous version of Backup Exec or Backup Exec for NetWare, except where end-of-life decisions have been made.

When you upgrade from previous versions, Backup Exec automatically converts your existing definitions, configurations, and jobs to the current version and some of your jobs may be combined or moved. After the migration completes, Backup Exec displays the Migration Report that you must review and acknowledge before Backup Exec continues the upgrade process. In this report, you can see how your jobs were migrated.

See [“Using the Migration Report to determine how existing jobs changed after an upgrade from a previous version of Backup Exec”](#) on page 109.

Backup Exec provides backward compatibility as follows:

- Backup Exec can communicate with earlier versions of the Remote Agent for Windows Systems.
You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.
- Backup Exec supports side-by-side installations of the Remote Administration Console.
The version of the Backup Exec server being remotely administered must be same as the version of the Remote Administration Console.

- If your environment includes a Central Administration Server (CAS) and Managed Backup Exec Servers (MBES), the CAS can only communicate with MBES that are running the same Backup Exec version or one version earlier. For example, if the CAS is running Backup Exec version N, CAS can only communicate with MBES running version N or N-1.

Before you upgrade Backup Exec, do the following:

- Delete the job histories and the catalogs that you no longer need to shorten the upgrade window.
- Run a database maintenance job.
- Verify that all available updates are installed for your current version of Backup Exec.
- Locate your license information and verify that your licenses are current. You must enter license information for Backup Exec when you upgrade. When you upgrade, on the **Add Licenses** panel, the existing licenses are listed. You can also add new entitlement IDs. Click **Next** and enter the credentials of the Veritas User Account that has permissions to the entitlements. Backup Exec automatically downloads the new license files from the Licensing Portal. If you already have the license file, use the **Import License File** button to import the license files.

You cannot change the configuration of your Backup Exec servers or the database location during the upgrade process. For example, you cannot change an administration server to a managed server. If you want to change the configuration of your Backup Exec servers, do it either before or after you upgrade to the current version. If you want to change the database location after the upgrade, use BEUtility.

You cannot change the directory where Backup Exec catalog is created during the upgrade process. It is also applicable for upgrade during silent mode installations.

Note: If you upgrade from a previous version of Backup Exec that uses a non-English version of Windows, you must download the SQL Server Express setup file for that language from the Microsoft website.

See [“About using a self-protect backup job to protect server metadata”](#) on page 230.

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 48.

Transport Layer Security (TLS) 1.2 protocol

Backup Exec only uses the Transport Layer Security (TLS) 1.2 protocol for secure communications with Backup Exec remote agents because TLS 1.0 has reached End-of-Life.

It is recommended that you upgrade the remote agents immediately after the Backup Exec server upgrade for secure communication.

Upgrade checklist for Backup Exec

Before you upgrade from a previous version of Backup Exec to the current version, do the following:

- Ensure that your backups are up to date. It is recommended that you always run full backups before and after you upgrade the applications or operating systems on any backup sources.
- Disable your antivirus software.
- Check the Backup Exec Software Compatibility List (SCL) and the Hardware Compatibility List (HCL) to verify that the applications that you want to back up and your storage devices are supported with this version of Backup Exec.
- Download all available upgrades and hot fixes for the version of Backup Exec that you want to install.
- Plan to perform the upgrade when system downtime won't affect users.
- Ensure that your entitlement IDs or License Files are available. You must enter new Backup Exec license information during the upgrade.
For details on updated licensing, see the *Backup Exec Licensing Guide*
- Review the following topics in the *Backup Exec Administrator's Guide*:
 - *System requirements for Backup Exec*
See [“System requirements for Backup Exec”](#) on page 45.
 - *How data lifecycle management (DLM) deletes expired backup sets on disk-based storage*
See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.
 - *About upgrading from previous versions of Backup Exec*
See [“About upgrading to Backup Exec”](#) on page 106.
- Review the document *Best practices for installing Backup Exec* on the Backup Exec knowledge base.

Using the Migration Report to determine how existing jobs changed after an upgrade from a previous version of Backup Exec

When you upgrade to Backup Exec, your existing definitions, configurations, and jobs are converted automatically to the current version. When the jobs are migrated, some of the jobs may be combined or moved. After the migration completes, Backup Exec displays the Migration Report for you to view and acknowledge. In this report, you can see how your jobs were migrated. The information that is included in the Migration Report cannot be recreated after the upgrade completes.

It is recommended that you review the Migration Report thoroughly to determine how your existing jobs have changed and how you may need to adjust your jobs manually. The Migration Report is available for viewing from the Backup Exec Administration Console or the **Home** tab at any time after the migration completes.

The Migration Report is stored in the following path:

```
<Backup Exec install path>\Backup
Exec\Data\MigrationReportFiles\Data-Migration-Report.html.
```

Note: To view the Migration Report, JavaScript must be enabled in your web browser. If the server on which you installed Backup Exec does not have a browser with JavaScript enabled, you can copy the Migration Report to another server that has JavaScript enabled.

To view the Migration Report

- ◆ Do one of the following:
- To view the Migration Report from the Administration Console

Click the Backup Exec button, select **Installation and Licensing**, and then select **Migration Report**.
- To view the Migration Report from the Home tab

On the **Home** tab, in the **Installation and Upgrades** group, click **Migration Report**.

See [“Configuring the Home tab”](#) on page 120.

Post-installation tasks

For best results, do the following after installing Backup Exec:

- Create disk storage so that Backup Exec can automatically manage the lifecycle of your backup data.

See [“Configuring disk storage”](#) on page 334.

- Make sure that your storage devices are connected and configured properly.
- Decide what types of storage devices you want to use for your backup jobs. You can configure storage devices when you prepare your Backup Exec environment.
- Understand how Backup Exec provides overwrite protection for your tape media.
See [“Media overwrite protection levels for tape media”](#) on page 505.
- Understand the default media set for tape media and its four-week overwrite protection period.
See [“Default media sets”](#) on page 491.
- Understand Data Lifecycle Management.
See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.
- Learn about creating new media sets with different retention periods.
See [“Creating media sets for tapes”](#) on page 499.
- Decide which credentials you want your Backup Exec logon account to use when browsing and making backup selections. You can use an existing Backup Exec logon account, or create a new one.
See [“Backup Exec logon accounts”](#) on page 759.
- Configure a secure connection to the Backup Exec Database
See [“Configuring encryption for the connection to the Backup Exec Database”](#) on page 708.

Uninstalling Backup Exec

Use Microsoft's Add or Remove Programs feature to remove Backup Exec from a computer. For additional information on Add or Remove Programs, refer to your Microsoft documentation.

Uninstalling Backup Exec also removes tape class drivers. If you reinstall Backup Exec and want to use tape class drivers, you must reinstall them.

To uninstall Backup Exec

- 1 Close Backup Exec.
- 2 From the Windows Control Panel, select the option to uninstall a program.
- 3 Select **Backup Exec™**, and then click **Uninstall**.

- 4 When you are prompted to confirm that you want to uninstall Backup Exec from your computer, click **Yes**.
 - 5 Select whether you want to remove only the Backup Exec program files or Backup Exec and all of its associated files.
 - 6 Click **Next**.
If the uninstall program fails, click **View Installation Log File** for additional information.
 - 7 If you are prompted, restart the computer.
- See [“Uninstalling Backup Exec features from the local Backup Exec server”](#) on page 111.

Uninstalling Backup Exec using the command line

If Backup Exec is already installed, you can use the setup.exe program to uninstall Backup Exec program files and Backup Exec data.

To uninstall Backup Exec using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change to the following directory:
`\be\winnt\install\bex64`
- 4 To remove the Backup Exec program files but keep all of the Backup Exec data, type:

```
SETUP /UNINSTALL:
```

To remove the Backup Exec program files and the Backup Exec data, type:

```
SETUP /REMOVEALL:
```

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 90.

Uninstalling Backup Exec features from the local Backup Exec server

The Installation Wizard removes Backup Exec features from the local Backup Exec server. All corresponding files, registry keys, and configurations are removed.

Note: License files remain on the server after features are uninstalled. Do not delete the license files while Backup Exec is installed. Deleting the license files causes the trial version to go into effect.

To uninstall Backup Exec features from the local Backup Exec server

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Features and licenses on this Backup Exec Server**.
- 2 On the **Add Licenses** panel, click **Next**.
- 3 On the **Configure Features** panel, deselect the agents or features that you want to uninstall, then click **Next**.
- 4 Uncheck the check box for the language that you want to remove.
- 5 On the **Choose Languages** panel, click **Next**.
- 6 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 7 Read the installation summary, and then click **Install** to start the process.
- 8 When the Installation Wizard has completed, click **Finish**.

See [“Uninstalling Backup Exec”](#) on page 110.

Getting Started


This chapter includes the following topics:

- [About the Backup Exec Administration Console](#)
- [Displaying the version information for Backup Exec](#)
- [Locking and unlocking the Backup Exec Console](#)
- [How to sort, filter, and copy information on the Backup Exec Administration Console](#)
- [Customizing views on the Backup Exec Administration Console](#)
- [Configuring the Home tab](#)
- [Configuring the RSS Reader](#)
- [Creating a disaster preparation plan \(DPP\)](#)
- [Getting started with backups](#)

About the Backup Exec Administration Console

Use the Backup Exec Administration Console to run backups, restore data, monitor jobs, configure storage, and run reports.

Figure 3-1 Backup Exec Administration Console



The administration console screen includes the following components:

Table 3-1 Administration console components

Label	Item	Description
1	Backup Exec button	The Backup Exec button displays on the upper left side of the administration console. To display the options in the Backup Exec button, click the Backup Exec button, select the menu name, and then select an option. You can launch Backup Exec operations by clicking options from a menu.

Table 3-1 Administration console components (*continued*)

Label	Item	Description
2	Tabs	Tabs at the top of the screen let you navigate Backup Exec. You can access the following views from the navigation bar: ■ Home. Provides quick access to the Backup Exec information that you use frequently. Customize the Home view by adding or deleting items. ■ Backup and Restore. Create a backup or restore job. ■ Job Monitor. Monitor and manage backup, restore, installation, and storage operation jobs. ■ Storage. Configure storage, run storage operations, and manage media. ■ Reports. View, print, save, and schedule reports about the Backup Exec server, operations, and device and media usage, and create custom reports. View reports in Backup Exec in PDF or HTML format, and save and print reports in PDF, XML, HTML, Microsoft Excel (XLS), and comma-separated value (CSV) formats
3	Groups	Groups display on the tabs in the administration console, and contain the commands that initiate actions such as creating a new backup job or configuring storage. The commands in the groups are dynamic, changing according to the selection. Some commands are unavailable until you select an item on the console screen or until you run a prerequisite task.
4	Columns	Customize columns by doing any of the following actions: ■ Drag and drop columns to change their location. ■ Right-click a column heading to select the columns to display or to sort and filter the column content. ■ Click the column heading to change the order of the column. For example, names of reports display in alphabetical order by default. To display report names in reverse alphabetical order, click the Name column heading on the Reports view.
5	Selection pane	Select items to work with, such as servers to back up or restore.

Table 3-1 Administration console components (continued)

Label	Item	Description
6	Details pane	Additional details appear for the server that you select in the list of servers. The Details pane lists the resources for the selected server and the selection status, backup status, and logon account information for each resource. When you select a Hyper-V host or VMware host in the list of servers, the Details pane includes the following details: ■ The last 7 days of backup jobs ■ The date of the last backup ■ The date of the next scheduled backup You can also restore data and filter the list of guest virtual machines from this pane.
7	Status bar	The status bar appears on the bottom of the administration console and provides information about the Backup Exec server, jobs that are running or are scheduled, alerts, and services.
	Refresh	Click F5 to refresh the user interface on the administration console.

Displaying the version information for Backup Exec

You can display information about the version of Backup Exec that is installed.

To display the version information for Backup Exec

- 1 Click the Backup Exec button, select **Help and Documentation**, and then click **About Backup Exec**.
- 2 Click **OK**.

Locking and unlocking the Backup Exec Console

To lock the Backup Exec session that you are working on and secure the Backup Exec console from unauthorized access, you can use the **Lock Console** feature.

You can only enable this feature, when you select the **Secure the Backup Exec console** check box in the **Network and Security** settings. The **Lock Console** feature is now enabled.

If you do not select the **Secure the Backup Exec console** check box, the **Lock Console** feature is grayed out.

See [“Changing network and security options for Backup Exec”](#) on page 718.

After you lock the Backup Exec user interface, you must enter the password to connect to the Backup Exec console. Unless you unlock the Backup Exec console, you cannot perform any tasks in the Backup Exec user interface.

Backup Exec has other utilities that can be launched from the console and have separate user interfaces. If any of these utilities are open when you lock the console, you can continue to use them. For example, Quick Assist, Install Options and Licenses, Help, Backup Exec Services, Install Agents, create Disaster Recovery Disk, and so on.

In case of rolling upgrade, if you have an earlier version of MMS and an updated version of CAS, and you connect to MMS from CAS, this feature is available but in disabled state (grayed out).

To lock and unlock the Backup Exec console

- 1 Click the Backup Exec button and select **Configuration and Settings > Backup Exec Settings > Network and Security**.
- 2 Select the **Secure the Backup Exec console** check box.

Note: By default, this check box is not selected.

- 3 Click the Backup Exec button, and then select **Lock Console**.
The **Connect to Backup Exec server** dialog box is displayed. The server name and user name is disabled.
- 4 Enter the password for the user name that you used to log on to the Backup Exec console.
- 5 Click **Connect**.
The state of the tasks before you locked the Backup Exec console resumes. You can now continue using Backup Exec.

How to sort, filter, and copy information on the Backup Exec Administration Console

You can customize the information that displays on the **Backup and Restore** tab, the **Job Monitor** tab, and the **Storage** tab.

You can do any or all of the following actions:

- Choose a default configuration that Backup Exec provides, such as **Servers with Active Jobs** or **Failed Jobs**.
- Specify a sort order for the columns that appear in the views.
- Specify the values that you want to use to filter the information that Backup Exec displays.
- Specify the columns that you want to appear and the order in which they should appear.
- Create and save a configuration to use again.
- Copy list items to the Clipboard and then copy them to any application that supports copy-and-paste.

To sort or filter information on the Backup Exec Administration Console

1 Do one of the following:

To customize a view of the computers in the list of servers	On the Backup and Restore tab, in the Views group, click Sort and Filter .
To customize a view of the storage devices	On the Storage tab, in the Views Group, click List , and then click Sort and Filter .

2 Do any of the following:

To select a default configuration, such as Servers with Failed Backups , or to select a configuration that you created and saved previously	Click Configurations and select a configuration.
To specify an ascending or descending sort order for the columns	Click Sort , choose the options as appropriate, and then click OK .
To specify one or more columns to filter for specific values	Click Filter , choose the options as appropriate, and then click OK .
To specify the columns that you want to display and the order in which they should appear	Click Columns , choose the options as appropriate, and then click OK .
To create and save a configuration	Click Save , choose the options as appropriate, and then click OK .

To edit a configured view

- 1 Do one of the following:

To edit a configuration from the Backup and Restore tab	On the Backup and Restore tab, in the Views group, click Sort and Filter .

To edit a configuration from the Storage tab	On the Storage tab, in the Views group, click Sort and Filter .

- 2 Click **Configurations**.

- 3 Select the configuration that you want to edit, and then click the pencil icon.

To delete a configured view

- 1 Do one of the following:

To delete a configuration from the Backup and Restore tab	On the Backup and Restore tab, in the Views group, click Sort and Filter .

To delete a configuration from the Storage tab	On the Storage tab, in the Views group, click Sort and Filter .

- 2 Click **Configurations**.

- 3 Select the configuration that you want to delete, and then click the delete icon.

To copy information on the Backup Exec Administration Console

- 1 On any tab except the **Home** tab, right-click an item in the list view.

- 2 Click **Copy**.

- 3 Open any application that supports copy-and-paste, and then paste the information.

Customizing views on the Backup Exec Administration Console

You can customize how you view the information that displays on the **Backup and Restore** tab, the **Job Monitor** tab, and the **Storage** tab.

Table 3-2 Views on the Backup Exec Administration Console

View	Description
Standard	Displays the information in a view that provides descriptive text.
Compact	Displays the information in a view that conserves space.
Tree	Displays the items in a hierarchical view. This view is disabled for the list of servers on the Backup and Restore tab.
List	Displays the items in a list that you can sort by columns. This view is disabled for the list of servers on the Backup and Restore tab.

To customize views on the Backup Exec Administration Console

- ◆ On the **Backup and Restore** tab, the **Job Monitor** tab, or the **Storage** tab, in the **Views** group, click **Standard**, **Compact**, **Tree**, or **List**.

Note: The **Tree** and **List** view is disabled for the list of servers on the **Backup and Restore** tab.

Configuring the Home tab

You can configure the **Home** tab by selecting the items that you want to display. You can drag and drop items to move them to another location on the **Home** tab or you can maximize a single item. The **Home** tab items contain Backup Exec data and links to features that you use frequently.

To configure the Home tab

- 1 On the **Home** tab, in the **Layout** group, click the layout for the items that you want to display.
- 2 In the **System Health** and **Support** groups, select the check box for the items that you want to display.
- 3 Drag the items to a column and position in which you want them to display to further customize the **Home** tab.

You can configure the **Layout**, and hide or display items in the **System Health** and the **Support** groups.

Restoring the Home tab's default layout

You can quickly restore the **Home** tab to its default configuration at any time.

To restore the Home tab's default configuration

- ◆ On the **Home** tab, in the **Layout** group, click **Reset Home Tab**.

Layout group

You can select one of the following layout configurations to display the items on the **Home** tab.

Table 3-3 Home Tab Layout items

Item	Description
One Column	Displays the Home tab items in one column.
Two Columns	Displays the Home tab items in two columns.
Narrow/Wide	Displays the Home tab items in two columns with a narrow panel and a wide panel.
Three Columns	Displays the Home tab items in three columns.
Reset Home Tab	Restores the contents of the Home tab to the default configuration.

System Health group

The items in the **System Health** group provide overviews of alerts, backup jobs, backup size data, storage status, and Simplified Disaster Recovery. You can select the following items to display on the Backup Exec **Home** tab.

Table 3-4 System Health group items

Item	Description
Active Alerts	Lets you view all alerts that have not received a response. You can filter the alerts to view specific types of alerts, the source of the alerts, and the amount of time that alerts occurred. You can display any or all of the following types of alerts: ■ Error ■ Warning ■ Attention Required ■ Information

Table 3-4 System Health group items (*continued*)

Item	Description
Alert History	Lets you view the property and response information for alerts.
Backup Status	Provides a summary view of the backup job status for the servers that are backed up or available for backup.
Backup Size	Provides a summary view of the amount of data that is backed up. You can customize the number of days for which you display information about the backup size. You can also select the type of backups that display.
Storage Status	Provides a summary view of the amount of space that is available on your storage. The storage information includes the total capacity that displays the amount space that is used for the different types of data.
Simplified Disaster Recovery	Provides a status of whether the Simplified Disaster Recovery disk image has been created. If the Simplified Disaster Recovery disk image has not been created, you can click the Create Disaster Recovery link to launch a wizard that guides you through the process. See “About Simplified Disaster Recovery” on page 877.
Database Encryption Key	Provides a status for whether the database encryption key has been exported. If you have not exported the database encryption key, you should do so to ensure that you can access the Backup Exec Database later. You are required to provide the database encryption key for disaster recovery and migration scenarios, for example. See “Exporting the Backup Exec Database encryption key” on page 705.

Table 3-4 System Health group items (*continued*)

Item	Description
Virtual Machine Backups	Lists the number of virtual machines that are being backed up with the virtual-based backup method and with the agent-based backup method. The information is updated when a backup job is created, deleted, or edited. It is also updated when the user interface is refreshed or relaunched.
Licensed Status	Provides a summary of the licensed usage status on the Backup Exec server. For more information, refer to the Licensing guide.
Anomaly Detection	Provides a summary of the anomalies detected for a backup job. Anomaly count: Displays the anomaly count where no action is taken. You can select to view the anomaly detection count for the last 7 days, 15 days, or 30 days. See “About Anomaly Detection” on page 297.

Support group

The items in the **Support** group provide technical support, documentation, licensing contracts, and the RSS Reader resources. You can select the following items to display on the Backup Exec **Home** tab.

Table 3-5 Support group items

Item	Description
Technical Support	Provides the following support options to help you understand product features and functionality or troubleshoot issues: ■ Backup Exec Tech Center ■ Backup Exec Technical Support ■ Best Practices ■ Use MySupport to manage new or existing support cases ■ Backup Exec Remote Assistance ■ Register to receive notifications ■ Get Backup Exec updates
Documentation	Provides the following documentation options to help you understand product features and functionality or troubleshoot issues: ■ View Readme ■ View Help (HTML) ■ View Administrator's Guide (PDF)
Licensing Information	Provides the following licensing options to help you manage license contracts and licenses: ■ Synchronize Entitlements with Licensing portal System ■ View license information ■ View license contract information ■ Purchase or renew licenses ■ Verify that Entitlements are updated to reflect the license renewal on Licensing Portal. ■ Update the installed license keys. ■ View licensing process information
RSS Reader	Lets you view and add Backup Exec and RSS feeds.

Table 3-5 Support group items (*continued*)

Item	Description
Installation and Upgrades	Provides the following reports: ■ Migration report ■ Installation Summary report This item only displays when you upgrade Backup Exec to a later version. See “Using the Migration Report to determine how existing jobs changed after an upgrade from a previous version of Backup Exec” on page 109. See “Viewing the Installation Summary Report” on page 99.

Managing instance-based usage

Backup Exec is moving to Subscription-based licensing with the aim to simplify and consolidate all licenses into a single edition. This edition includes all features offered by Backup Exec including the new Microsoft 365 feature. The subscription metering is instance-based. An instance is a virtual machine or physical machine or 10 Microsoft 365 users.

The calculation is based on backups done in the past 30 days and you can manage your instance-based license usage. If you back up more number of instances than what you are entitled for, Backup Exec goes into a 30-day Grace Period state that is followed by the Enforcement state. In the Enforcement state, backup jobs for only the entitled instances are allowed to run. Backup jobs for the remaining instances fail with a licensing error. You can change your selection any number of times when Backup Exec is running in the Grace Period state and only three times when Backup Exec is running in Enforcement state.

For details on updated licensing, see the [Backup Exec Licensing Guide](#).

To manage instance-based usage

- 1 In the **System Health** group, under **Licensed Status**, click **Manage Instances**.

The **Manage Instance Usage** panel displays the following fields:

Physical/Virtual Machines	The number of physical or virtual machines that are protected. Each physical or virtual machine is equal to one instance.
Microsoft 365 Users	The number of instances in which data of Microsoft 365 users is protected. Each instance is equal to 10 Microsoft 365 users. For example, if you are protecting 8 Microsoft 365 users, it is displayed as 1 (8 users). Here you are protecting 1 instance that has 8 users. If you are protecting 20 Microsoft 365 users, it is displayed as 2 (20 users). Here you are protecting 2 instances that have 20 users.
Total selected instances	The number of instances that are currently selected.
Total entitled instances	The number of licensed instances.

- 2 Select the **Physical/Virtual Machines** or the **Microsoft 365 Users** tab.

The table within each tab displays the instances or Microsoft 365 users backed up in the last 30 days. You can also search for an instance name or user using the **Search** option.

- 3 If you select the **Physical/Virtual Machines** tab, you can view or update the table depending on whether you are in the grace period or enforcement state.

Instance Name	Server or machine name that is backed up.
VM Resource	Display name of the virtual machine, if applicable.
Clustered/Distributed Application Node	Server or machine name that is part of a cluster, SharePoint, or Enterprise Vault.
Job Name	Name of the last backup job that was run for the server.

- 4 If you select the **Microsoft 365 Users** tab, you can view or update the table depending on whether you are in the grace period or enforcement state.

User ID	Unique Microsoft ID of the user.
Tenant ID	Unique Microsoft ID of the tenant.
Application	Name of the Microsoft 365 application that is protected.
Job Name	Name of the backup job.
Backup Start Time	Time when the backup started.

- 5 (Optional) Click **Export** to create a CSV file of the instance usage details.
- 6 Click **OK**.

Managing capacity-based usage

This license includes all features offered by Backup Exec including the Microsoft 365 feature. It allows you to protect total number of front-end terabytes (FETB) of your data. The subscription metering is capacity-based.

The capacity is calculated for a virtual/physical machine and Microsoft 365 user entity. The calculation is based on largest full backup set for each resource backed up in the past 30 days.

If you back up more capacity than what you are entitled for, Backup Exec goes into a 30-day Grace Period state that is followed by the Enforcement state. In the Enforcement state, all the backup jobs start to fail.

With this license you can remove a maximum of 3 resources from the virtual/physical machine or Microsoft 365 backup to make changes to your capacity usage. You can do this only once.

For details on updated licensing, see the [Backup Exec Licensing Guide](#).

To manage capacity-based usage

- 1
- In the **System Heath** group, under **Licensed Status**, click **View details** for the Capacity-based license.

The **Capacity Usage Calculation Information** panel displays the following fields:

Physical/Virtual Machines	The resources that are backed up and size of the resource that is backed up for the physical/virtual machines.
Microsoft 365 Workloads	The entities that are backed up and the size of the entities that are backed up for the Microsoft 365 workloads.

- 2
- Select the **Physical/Virtual Machines** or the **Microsoft 365 Workloads** tab.

The table within each tab displays the backup sets and amount of data that is backed up in the last 30 days. You can also search for a backup source for a physical or virtual machine and by entity name for Microsoft 365 workloads using the **Search** option.

- 3
- If you select the **Physical/Virtual Machines** tab, you can view or update the table depending on whether you are in the grace period or enforcement state.

Backup Source	Resource or machine that is backed up.
Resource Type	Type of resource, whether it is a physical or virtual machine.
Backup Time	Date and time of the backup job that was run on the server.
Job Name	Name of the backup job that was run on the server.
Size	Size of the data that is backed up.

Right-click the backup set and click **Remove from Capacity Usage Calculation** to remove the resource from the physical or virtual machine to make changes to your capacity usage.

- 4 If you select the **Microsoft 365 Workloads** tab, you can view or update the table depending on whether you are in the grace period or enforcement state.

Entity Name	Name of the Microsoft 365 backup item or entity.
Backup Source	Microsoft 365 resource that is backed up.
Tenant Name	Name of the Microsoft 365 tenant.
Backup Time	Date and time of the backup job that was run on the server.
Job Name	Name of the backup job that was run on the server.
Size	Amount of data that is backed up.

Right-click the backup set and click **Remove from Capacity Usage Calculation** to remove the resource from Microsoft 365 backup to make changes to your capacity usage.

- 5 (Optional) Click **Export** to create a CSV file of the instance usage details.
- 6 Click **OK**.

Managing restore only license

The restore only license in Backup Exec allows you to restore your backup sets if you do not have a valid license or if your license has expired. If you have a valid backup license, your backups and restores continue to work as expected.

Every restore run is part of the usage calculation even if there are multiple restores for the same backup set. Based on the start and end date, while you are using the license, every time the restore runs, it is displayed on the Backup Exec user interface. When the license expires or when you reach the total capacity, the restore jobs fail.

To manage restore only license

- 1 In the **System Health** group, under **Licensed Status**, click **View details** for the Restore-based usage.

The **Restore Capacity Usage Calculation Information** pane displays the following fields:

Restore Time	Date and time when the restore run was initiated.
Job Name	Name of the restore job that was run on the server.
Size	Size of the data that is restored with each restore job.

- 2 (Optional) Click **Export** to create a CSV file of the instance usage details.
- 3 Click **OK**.

Configuring the RSS Reader

You can customize the RSS Reader and select the default Backup Exec feeds that display in the reader. You can add additional Backup Exec RSS feeds or remove RSS feeds.

The RSS Reader sorts articles by the date and the time. The reader displays the last entry of an article in the RSS feed; however, you can choose to view the full article.

The RSS Reader refreshes the RSS feeds every 15 minutes when the item is open in the **Home** tab. If the RSS feed is not open in the reader, the RSS feed does not refresh.

To view an article in the RSS Reader

- 1 On the **Home** tab, in the **Support** group, select the **RSS Reader** check box.
- 2 In the **RSS Reader**, click the arrow next to the RSS feed that contains the article.
- 3 Click the hyperlink for the article that you want to open.

The RSS Reader opens a new window that contains a portion of the article from the RSS feed.

- 4 Click **Go to full Article** to open Internet Explorer and view the entire contents of the article.

To customize the RSS feeds to the RSS Reader

- 1 On the **Home** tab, in the **Support** group, select the **RSS Reader** check box.
- 2 In the **RSS Reader**, click the pencil icon to add an RSS feed.
- 3 Enter any of the following:

URL	Indicates the location of the RSS feed that you want to add to the RSS Reader.
Name	Indicates the name of the RSS feed that you want to display in the RSS Reader.
Click here to see more RSS feeds	Shows a list of RSS feeds that you can add to the RSS Reader.

- 4 Click **OK**.

To remove an RSS feed from the RSS Reader

- 1 On the **Home** tab, in the **Support** group, select the **RSS Reader** check box.
- 2 Do one of the following:

To remove a default Backup Exec RSS feed	Clear the check box of the Backup Exec RSS feed.
To remove an RSS feed that you added to the RSS Reader	Click the red X next to the name of the RSS feed.

Creating a disaster preparation plan (DPP)

Disaster preparation planning is the implementation of strategies and procedures that minimize damage in the event a catastrophe destroys your data.

The following basic methods are available for disaster recovery:

- Automated recovery. Backup Exec's Simplified Disaster Recovery (SDR) feature automates the disaster recovery process for Windows computers.
- Manual recovery. You can manually recover both local and remote Windows computers.

The purpose of a Disaster Preparation Plan (DPP) is to return to an operational status as quickly as possible. Backup Exec is a crucial component of the DPP. The DPP you put in place with your Backup Exec system should be customized to your network environment.

For more information about how to perform a manual disaster recovery, refer to the following sections:

See [“Performing manual disaster recovery of a local Backup Exec server on a Windows computer”](#) on page 930.

See [“Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer”](#) on page 935.

While environments vary in different organizations, consider the following elements when creating a comprehensive DPP.

Table 3-6 Key elements of a DPP

Element	Description
Hardware protection	The hardware devices on your network (CPUs, drives, video) are susceptible to damage from many disaster situations. Uninterruptible power supplies (UPS), surge protectors, and security monitoring devices are the equipment most often used today to protect hardware. If you do not already have these items in place, you should consider installing them. The initial investment could be justified many times over in the event of a disaster.
The ability to maintain business operations during a disaster period	Make sure that proper precautions are taken by everyone to implement plans for network interruptions. For example, the phones in the sales department won't stop ringing because the server is down, so orders may have to be handwritten until the server is up again. Each department should work out strategies for such occurrences. If the proper precautions are taken, the server can be rebuilt quickly and operations can still continue.
A sound backup strategy.	A well-designed backup strategy that includes a strong media rotation scheme plays a key role in quickly restoring your file server.
Off-site and duplicate stage backups.	It is imperative that you regularly move the backed-up data to an off-site storage facility. If you use disk as your storage medium, consider adding a stage to duplicate backups to other storage. This ensures that if something happens to your facility, all of your backups are not destroyed. Depending on the importance of your data, you may choose to use several off-site storage facilities.

Table 3-6 Key elements of a DPP (*continued*)

Element	Description
Effective DPP management	A person or group of people should constantly supervise your organization's disaster preparation efforts. This person or group should install and maintain hardware protection devices, make sure all departments have a plan if the server goes down temporarily, and make sure that backups are made and rotated off-site regularly. Document your Disaster Preparation Plan for reference purposes.

See [“Getting started with backups”](#) on page 133.

Getting started with backups

After you install Backup Exec, you can run a backup job. The following table describes the recommended process for getting started.

Table 3-7 Getting started with backups

Step	Description
1. Configure storage	You must have a configured storage device before you can run any backup jobs. If no storage devices are already configured, such as tape drives or robotic libraries, then you can use the Configure Storage wizard to set up storage devices in Backup Exec. The wizard guides you through configuring all of the storage that Backup Exec supports. To start the wizard, on the Storage tab, in the Configure group, click Configure Storage. The wizard guides you through the rest of the process. See “Using the Configure Storage wizard” on page 538.

Table 3-7 Getting started with backups (*continued*)

Step	Description
2. Configure Backup Exec logon accounts	You can use the default Backup Exec logon account, which is the system logon account for the Backup Exec server. Backup Exec uses the system logon account by default for most backups. The system logon account contains the credentials of the Backup Exec service account. If the service account does not have rights to access the data that you select for backup, you can use the Logon Account Wizard. Use this wizard to create additional logon accounts that do contain the necessary credentials for accessing that data. To start the Logon Account Wizard, click the Backup Exec button. Select Configuration and Settings > Logon Accounts > Logon Account Wizard. See “Backup Exec logon accounts” on page 759.
3. Run a backup job	You can back up any of the computers on which you installed the appropriate Backup Exec agent, such the Agent for Windows. If you didn't install a Backup Exec agent on a computer, you can do so now. On the Backup and Restore tab, in the Servers group, click Add, and then click the appropriate selection. The wizard guides you through the rest of the process. To start a backup job, on the Backup and Restore tab, select the computer that you want to back up and then in the Backups group, click Backup. Click a menu item, such as Back Up to Disk. You can click Edit to change any of the defaults, or you can click OK to accept all of the defaults and let Backup Exec schedule the job. See “Backing up data” on page 158.

Table 3-7 Getting started with backups (continued)

Step	Description
4. Create the Simplified Disaster Recovery disk image	By default, Backup Exec selects all the data on a computer for backup, including the critical system components that you need to perform a full system restore by using SDR. SDR-enabled backups are those backups for which all of the critical system components are selected for backup. You must have SDR-enabled backups to use Backup Exec to rebuild a computer and restore it to a functional state. See “How to ensure that backups are enabled for Simplified Disaster Recovery” on page 885. After you have run SDR-enabled backups for all of the computers that you want to protect, you should create a Simplified Disaster Recovery (SDR) disk image. On the Backup and Restore tab, in the Restores group, click Create Disaster Recovery Disk. The wizard guides you through the rest of the process. See “Creating a Simplified Disaster Recovery disk image” on page 892.

Backups

This chapter includes the following topics:

- [How to prepare for your first backup with Backup Exec](#)
- [Recommendations for when to use virtual-based backup and agent-based backup](#)
- [Improving backup performance in Backup Exec](#)
- [Troubleshooting backup performance](#)
- [Required user rights for backup jobs](#)
- [About the list of servers on the Backup and Restore tab](#)
- [Adding servers that you want to back up to the list of servers on the Backup and Restore tab](#)
- [Removing servers from the list of servers on the Backup and Restore tab](#)
- [Creating a server group](#)
- [Hiding or viewing server groups on the Backup and Restore tab](#)
- [Adding servers to a server group](#)
- [Removing servers from a server group](#)
- [Editing a server group](#)
- [Moving servers to the Retired Servers server group](#)
- [Moving retired servers back to the All Servers server group](#)
- [Removing server groups from Backup Exec](#)
- [Backing up data](#)

- About selecting data to back up
- Changing the order in which backup sources are backed up
- Excluding files from backups
- Including specific files with a backup job's backup selections
- About backing up critical system components
- Backup Exec Shadow Copy Components file system
- Backup methods in Backup Exec
- Configuring backup methods for backup jobs
- How Backup Exec determines if a file has been backed up
- Configuring Backup Exec to automatically delete files after a backup
- Configuring network options for backup jobs
- Running the next scheduled backup job before its scheduled time
- Editing backup definitions
- Testing credentials for backup sources
- Replacing the credentials for a backup source
- Creating new credentials for a backup source
- Deleting retired or unused backup sources from the Credentials pane
- How job scheduling works in Backup Exec
- Including a specific date in the schedule for a backup job
- Preventing backup jobs from running on a specific date
- Viewing all scheduled backup jobs on a calendar
- Adding a stage to a backup definition
- Editing a stage
- Duplicating backup sets or a job history manually
- Running a test run job manually
- Verifying backed up data manually

- [Copying data from a virtual tape library to a physical tape device using DirectCopy to tape](#)
- [About using a self-protect backup job to protect server metadata](#)

How to prepare for your first backup with Backup Exec

Before you back up data, you should develop a backup strategy that includes the backup method, frequency, and data retention methods that are appropriate for your organization. A backup strategy is the collection of procedures you implement as a solution for backing up your environment. You may have different strategies for different areas of the organization.

You may need to analyze your backup environment to determine the following:

- How much data needs to be backed up?
- How long will the backups take?
- How much storage is required?

If you encounter degraded performance in your backup or restore jobs, follow the troubleshooting steps in the following sections to identify and resolve the issues.

See [“Improving backup performance in Backup Exec”](#) on page 142.

See [“Troubleshooting backup performance”](#) on page 145.

Before you run a backup job, you should ensure that you have the proper user rights.

See [“Required user rights for backup jobs”](#) on page 150.

You must also configure storage before creating backup jobs. You can set up Backup Exec to use specific storage devices or logical groupings of devices, such as storage pools.

Specifically, you might want to perform the following tasks to help you manage storage hardware and media most effectively:

- Create disk-based storage so that Backup Exec can automatically manage backup data retention.
See [“Configuring disk storage”](#) on page 334.
- Set up storage device pools to load-balance jobs.
See [“Creating storage device pools”](#) on page 527.
- Create media sets to manage data retention for tape cartridge media.

See [“Default media sets”](#) on page 491.

- Configure deduplication disk storage to optimize storage and network bandwidth.

See [“About the Deduplication feature”](#) on page 969.

Recommendations for when to use virtual-based backup and agent-based backup

Backup Exec offers two backup methods for virtual machines: virtual-based backup and agent-based backup.

Virtual-based backup

- Traditional

In the traditional method, all full, incremental, and differential backups are taken from the source virtual machine.

- Forever Incremental

In Forever Incremental backups a full backup is taken from the source virtual machine and then followed by incremental backups. A consolidate backup is run by consolidating the previous set of full and incremental backups.

Subsequent incremental backups use the consolidate full backup as a baseline to determine changes in the source virtual machine. A consolidate full backup is equivalent to a full backup from the source virtual machine. If you select deduplication disk storage for forever incremental backups, the Backup Exec Deduplication Block Cloning technology is used that significantly improves backup performance.

See [“Differences between the traditional and the forever incremental backups of virtual machines”](#) on page 141.

See [“About Forever Incremental Backup”](#) on page 940.

A virtual-based backup requires the Agent for VMware and Hyper-V to be installed on the Backup Exec server. For Hyper-V environments, the Agent for VMware and Hyper-V must also be installed on the Hyper-V host. In addition, the virtual machine's host or the vCenter Server must be added to Backup Exec as a server. This backup method uses VMware VADP APIs or Microsoft VSS snapshots to protect virtual machines.

Virtual-based backup is recommended for the following:

- Backups that require file/folder-level Granular Recovery Technology (GRT)

Note: The Agent for Windows must be installed on the virtual machine to restore file/folder-level GRT items back to their original location.

- Single-server, application-level GRT for Active Directory, Exchange, SQL, and SharePoint.

Note: The Agent for Windows must be installed on the virtual machine to perform application-level GRT.

- Backups that do not require GRT
- Offline virtual machines
- Disaster recovery
- Virtual machine templates.

Virtual-based backup is not recommended for the following:

- Backups of Oracle, Enterprise Vault, distributed SharePoint farms, and Exchange DAG. Virtual-based backup does not support application-level GRT for those items.

Agent-based backup

An agent-based backup requires the Agent for Windows to be installed on the guest virtual machine. This backup method is sometimes referred to as legacy backup, traditional backup, or in-guest backup. With this backup method, a virtual machine is treated as a physical server.

Agent-based backup is recommended for the following:

- Exchange DAG
- Oracle
- Enterprise Vault
- Distributed SharePoint farm
- A subset of files on the virtual machine (no system state)
- VMware Fault Tolerant virtual machines
- Virtual machines with Physical Raw Device Mapping (RDMS)

Agent-based backup is not recommended for the following:

- Disaster recovery of a virtual machine.
- Offline virtual machines

■ Virtual machine templates

Differences between the traditional and the forever incremental backups of virtual machines

The following table describes the differences between the traditional and the forever incremental backups of virtual machines.

Table 4-1 Differences between the traditional and the forever incremental backups of virtual machines

Item	Traditional backup	Forever Incremental backup
Supported backup types	Supports full, incremental, and differential backups.	Supports full, consolidate full, and incremental backups.
Connection to the source virtual machine	All backups from the source virtual machine. Connection required.	Only incremental backups from the source virtual machine. No connection to the source virtual machine is required for Consolidate Full backups.
Load on the source virtual machine	Increases load as all backups are from the source virtual machine.	Reduces the load as only incremental backups are from the source virtual machine.
Backup method Only applicable for Hyper-V virtual machines	Standard processing, Faster processing, and Resilient Change Tracking (RCT) methods.	Standard processing and Resilient Change Tracking (RCT) methods only.
Supported storage	All types of supported storage	Two types of storage are supported. ■ Disk storage and disk storage that is hosted on a network share ■ Deduplication disk storage
Block Cloning technology	Not applicable	Backup Exec Deduplication Block Cloning technology is used when backups are targeted to Deduplication disk storage.
Storage space consumed	High	■ Disk storage: High ■ Deduplication disk storage: Low

Table 4-1 Differences between the traditional and the forever incremental backups of virtual machines *(continued)*

Item	Traditional backup	Forever Incremental backup
Backup process time	High	■ Disk storage: High ■ Deduplication disk storage: Low

For differences in the backup operations, refer to the following section:

See [“Supported storage in forever incremental backups”](#) on page 946.

See [“About Forever Incremental Backup”](#) on page 940.

Improving backup performance in Backup Exec

Backup operations run in a group of systems. These systems can be compared to pipelines of various sizes, from the disk containing the data all the way to the backup destination. If any of these pipelines are constricted, they can become bottlenecks that cause the entire backup process to slow down. The troubleshooting steps in this section can help you identify bottlenecks in your backup or restore operations.

Some variables which can affect backup or restore performance include:

Item	Description
Hardware	Some hardware-related variables which can affect performance include: ■ Speed of the disk controller ■ Improper cabling or termination ■ Hardware errors that are caused by the disk drive, tape drive, disk controller, or SCSI bus Confirm that the controller is rated for the tape backup hardware. If it is not, you may experience unexpected performance limitations. Confirm that the SCSI BIOS Settings are set as follows: ■ Initiate Wide Negotiation is set to Yes when the tape device is connected to a 68-pin wide SCSI Cable Connector ■ Tape drives are not connected to a SCSI Raid Controller

Item	Description
System	The capacity and speed of the media server performing the backup, or the remote system being backed up, significantly affect performance. System activity during the backup also affects performance. Fragmented disks take a longer time to back up. Heavily fragmented hard disks not only affect the rate at which data is written to tape, but also affect overall system performance. Fragmented files take longer to back up because each segment of data is located at a different location on the disk. This fragmentation increases the time that is required to access the data. Make sure to defragment disks on a regular basis.
Memory	The amount of available memory affects backup speed. Insufficient memory, improper page file settings, or a lack of available free disk space can all cause excessive paging and slow performance. Make sure that every program and every process release the memory that it allocates when it starts. If a program or process does not release as much memory as it originally allocated, a memory leak occurs.
File types	An average file compresses at a 2:1 ratio when it is compressed using hardware compression. Higher and lower compression ratios occur depending on the type of files being backed up. If no compression is used, the tape device performs at its rated speed. Average compression ratios can double the backup speed. Image and picture files are fully compressed on disk. Hardware compression is performed by the tape device and not the backup software.
Compression	Successful compression can increase the tape drive's data transfer rate up to twice the native rate. Compression can be highly variable depending on the input data. Image files from a graphical program might compress at 4.5:1 or more, while binary files might compress at only 1.5:1. Data that has already been compressed or encrypted may expand by about five percent if you attempt to compress it further. This data expansion can reduce drive throughput.
Files	The total number of files on a disk and the relative size of each file affects backup performance. Disks containing fewer files, but where each file is large, run backups quickly. Backups run slower if a disk contains many small files. A large number of files which reside in the same directory path can be backed up more efficiently than files in multiple directory locations.

Item	Description
Block size	Larger block sizes can improve the compression ratio, which helps the drive achieve better throughput and more tape capacity. Make sure that the block sizes and buffer sizes are set properly. Throughput increases in proportion to the compression achieved, until you reach the drive's maximum throughput. Do not increase the block size beyond the default settings.
Network	The physical connection to a remote disk limits that disk's backup speed. The rate at which a remote server's hard disks can be backed up depends on the following factors: ■ The make and model of network cards. ■ The mode and frame type configuration for the adapter. ■ The connectivity equipment. ■ Windows Settings. ■ Location of the drives. Local disk drives on the media server can usually be backed up faster than remote servers across a network. A common cause of slow network backups is networking configuration. Features such as "full-duplex" and "auto-detect" may not be fully supported in every environment. To improve throughput, manually set the speed to 100 MB and the duplex to half/full for the server side. Find out which Ethernet port the server is connected to on the switch, and set the SWITCH PORT setting to 100 MB and half/full duplex. Do this for the backup server switch port, and any switch ports for computers being backed up. Note: When a hub is in place instead of a switch, full duplex may not be supported. See the Original Equipment Manufacturer for details on device features. Note: Both the switch and the network card must have matching settings. For instance, if the switch port is set to 100 half, the NIC for the server should also be set to 100 half. If a full duplex backup is slower than a half duplex backup, full duplex may not be supported for your combination of NIC, driver, and switch. Contact the NIC and switch manufacturer for updated drivers, firmware, or other support documentation. The NIC driver can be a common cause of slow throughput. The NIC driver can be easily overwritten by an operating system service pack. If a service pack has been applied and the driver has been overwritten, reinstall the Original Equipment Manufacturer driver.

Item	Description
Debugging	Debugging that is enabled for troubleshooting purposes can also affect system performance. Debugging that occurs through the Services applet is temporary. Cycle the services or restart the computer to stop the debugging. Debugging configured through the Windows Registry allows for continuous debugging. Leaving the services in debugging mode causes the logs to accumulate. To improve performance, either take the services out of debugging mode when the problem is resolved, delete the older debug files, or compress the logs directory.
Backup Exec Database	Installing the Backup Exec Database (BEDB) to an existing SQL instance that other applications use can also degrade performance. This is particularly relevant in a Central Administration Server (CAS) environment. Other applications may cause resource issues and use all the available resources within the instance.

Troubleshooting backup performance

You can perform several steps to identify the causes of any Backup Exec job performance issues that you encounter. This section examines performance troubleshooting for the following job types:

- Local backup to disk
- Remote Backup to disk
- Local backup to tape
- Remote Backup to tape

Local backup to disk

1. Get a baseline. Review previous jobs in the **Job History** window of the **Job Monitor** tab. Note both the speed of previous jobs and the overall duration of these backups. Observe the total time that jobs take to complete, rather than the actual byte count rates. If current jobs take longer to complete than previous jobs, or do not meet speed expectations, continue troubleshooting.
2. Narrow down the problem. If the backup job includes multiple drives or agents, split the job up into separate jobs for each of those drives and agents. You can then review the performance of each drive or agent separately. To split up a backup job, click on the Backup Exec button, select the **C\$** drive, schedule the job, and click **Submit**. If performance is slow only for a particular job, continue troubleshooting that job.

3. If a particular job still shows slow performance, split the job again to further determine which part of the data most affects the job's performance.

A section of data with many small files and directories will negatively impact performance. This performance impact is normal behavior.

Check whether the data is redirected somewhere else. Some file systems allow a directory to remotely mount data. The files in these directories can be located on remote servers, which may degrade performance for the entire backup.

4. Test backup-to-disk (B2D) throughput. Use Windows to copy at least 2 GB of data in the backup job to the B2D disk. Compare the performance of the Windows copy to the performance of the backup. If the performances of both are comparable, the performance bottleneck is likely in the disk subsystem in which the B2D folders reside. Move the B2D folders to a faster disk subsystem, or continue troubleshooting.
5. Test system throughput. If your job is file-based, instead of based on an Exchange, SQL, or other type of database backup, create a similar backup in NTBackup (Windows Backup) and perform a backup to disk. Compare the performance of the NTBackup job to the performance of the Backup Exec job.

If you need to back up an Exchange, SQL, or other database, create a backup-to-disk job in Backup Exec that backs up 2 GB of data to wherever that database agent resides. Perform the same test with NTBackup. Compare the performance of both backups. If performance rates are similar, then Backup Exec is performing at the capacity of the system.

Remote backup to disk

1. Get a baseline. Review previous jobs in the **Job History** window of the **Job Monitor** tab. Note both the speed of previous jobs and the overall time that is required for these backups. Observe the total time that jobs take to complete, rather than the actual byte count rates. If current jobs take longer to complete than previous jobs, or do not meet speed expectations, continue troubleshooting.
2. Narrow down the problem. If your job includes multiple drives or agents, split the job up into separate jobs for each of those drives and agents. You can then review the performance of each drive or agent separately. To split up a backup job, click on the Backup Exec button, select the **C\$** drive, schedule the job, and click **Submit**. If performance is slow only for a particular job, continue troubleshooting that job.
3. If a particular job still shows slow performance, split the job again to further determine if any particular part of the data affects performance. A section

of data with many small files and directories negatively affects performance. This performance impact is normal behavior.

Check whether the data is redirected somewhere else. Some file systems allow a directory to remotely mount data. The files in these directories can be located on remote servers, which may degrade performance for the entire backup.

4. Test network throughput. Copy between 500 MB and 1 GB of data from the backup server to the remote server. Note how long the copy operation takes to complete. You can perform this copy by creating a path to another server. In the Windows command line, type `<\\remote servername\c$>`. When the drive is displayed, copy the data.

Follow this same procedure to copy data from the remote server to the backup server, and note how long the operation takes to complete.

Compare the speed of both of these network tests with Backup Exec's performance. If Backup Exec performance is slower than the file copy tests, the network is likely not the bottleneck.

If the network is not the bottleneck, consider performing the same test to a different remote server, or between two different servers. This can help determine if the performance issue is associated with the network in general, or a particular server on the network. If you do not find any network performance issues, continue to the next step.

5. Test system throughput. Try to back up the remote server with NTBackup (Windows backup). If the remote server is not visible in NTBackup, create a mapped drive to the server's drive and try to back up at least 2 GB of data. Compare the NTBackup logs to the Backup Exec logs and identify any performance differences.

Local backup to tape

1. Get a baseline. Review previous jobs in the **Job History** window of the **Job Monitor** tab. Note both the speed of previous jobs and the overall duration of these backups. Observe the total time that jobs take to complete, rather than the actual byte count rates. If current jobs take longer to complete than previous jobs, or do not meet speed expectations, continue troubleshooting.
2. Clear any temporary hardware glitches. Turn off the power to the server, tape drive, or tape library, then turn it on again. Turn off the backup server first, then the tape drive or library. Wait a few seconds, then turn on the tape drive or tape library. When the tape drive or tape library is ready, turn on the server. Run the backup job again and examine its performance. If the performance issue persists, continue troubleshooting.

3. Check the SCSI subsystem. Slow performance can be caused by the disk drive, the tape drive, the disk controller, the SCSI bus, or improper cabling or termination. Ensure that the following are true:
 - The controller is rated for the tape backup hardware.
 - The SCSI BIOS Settings are set properly.
 - **Initiate Wide Negotiation** is set to **Yes** when the tape device is connected to a 68-pin-wide SCSI Cable Connector.
 - Tape drives are not connected to a SCSI Raid Controller.

The performance of the verify operation shows the health of the SCSI subsystem. Because the verify operation only reads data and performs in-memory operations on the media server, the speed of the SCSI subsystem limits the verify operation's performance. You can examine the performance of the verify operation by examining the job logs of any jobs which include a verify operation. If the verify speeds are slow, the SCSI subsystem is likely the performance bottleneck.

4. Split the job up into smaller jobs to identify which agents or features might affect performance. If any of the smaller jobs show performance issues, continue troubleshooting that job.
5. If a particular job still shows slow performance, split the job again to further determine if any particular part of the data negatively affects performance. A section of data with many small files and directories negatively affects performance. This performance impact is normal behavior.

Determine whether the data is redirected somewhere else. Some file systems allow a directory to remotely mount data. The files in these directories can be located on remote servers, which may degrade performance for the entire backup.

6. Test system throughput. Try to back up the remote server with NTBackup (Windows backup). If the remote server is not visible in NTBackup, create a mapped drive to the server's drive and try to back up at least 2 GB of data. Compare the NTBackup logs to the Backup Exec logs and identify any performance differences.
7. Successful compression can increase the tape drive's data transfer rate to twice the native rate. Compression performance can be highly variable depending on the input data. Image files can compress at a ratio of 4.5:1 or more. Binary files may compress at only a 1.5:1 ratio. Data that has already been compressed or encrypted may expand by about five percent if you attempt to compress it further. This expansion can reduce drive throughput.

If hardware or software compression does not perform as expected, switch to the other type of compression. You can switch compression types by editing the backup job properties, clicking on **General** under **Settings**, then selecting a different type of compression under the **Compression Type** menu.

Remote backup to tape

1. Perform any of the troubleshooting steps for **Local backup to tape** given above. You can also perform the following steps.
2. Test network throughput. Copy between 500 MB and 1 GB of data from the backup server to the remote server and note the duration of the copy operation. To do this, create a path to another server. In the Windows command line, type `<\\remote_servername\c$>`. When the drive is displayed, copy the data.

Follow this same procedure to copy data from the remote server to the backup server, and note how long the operation takes to complete.

Compare the speed of both of these network tests with Backup Exec's performance. If Backup Exec performance is slower than the file copy tests, the network is likely not the bottleneck.

If the network is not the bottleneck, you can perform the same test to a different remote server, or between two different servers. This may help you determine if the performance issue is associated with the network in general, or with a particular server on the network. If you do not find any network performance issues, continue to the next step.

3. Test system throughput. Try to back up the remote server with NTBackup (Windows backup). If the remote server is not visible in NTBackup, create a mapped drive to the server's drive and try to back up at least 2 GB of data. Compare the NTBackup logs to the Backup Exec logs and identify any performance differences.

Note: If remote backups are not possible with NTBackup, open NTBackup locally on the remote server and run a local backup job there. Use Backup Exec to back up the same data to disk, and compare the performance of both backups. In most cases, Backup Exec jobs which back up to disk run faster than those that back up to tape.

Required user rights for backup jobs

To perform any backup operations, the following Windows user rights are required for the service account and any Backup Exec logon accounts:

- Act as part of the operating system.
- Back up files and directories.
- Restore files and directories.
- Manage auditing and security log.
- Logon as a batch job (only for Windows Vista and later).

For more information about user rights in Windows operating systems, see your Microsoft documentation.

See [“Changing the credentials for a service account”](#) on page 771.

See [“Backup Exec logon accounts”](#) on page 759.

About the list of servers on the Backup and Restore tab

You can view a list of servers on the **Backup and Restore** tab. The servers that display in the list include any servers that Backup Exec discovered during an upgrade, any servers that you manually add to Backup Exec, and any servers that Backup Exec discovers during a catalog operation. Servers must be added to the list so that you can select them for backup jobs.

You can also monitor server activity and job status from the list of servers. By default, Backup Exec displays a server's alerts, backup status, and a calendar of the last seven days of backup jobs. It also displays the date and time of the previous and upcoming scheduled backups. You can customize the columns on this list to display additional information.

You can select to view any of the following details about each server in the list:

- Server
- Server type
- Version of the server
- Backup Exec version
- Data source types
- Backup selections

Adding servers that you want to back up to the list of servers on the Backup and Restore tab

- Active alerts
- Status
- Last seven days of backup jobs
- Last backup
- Next backup
- Percent complete
- Elapsed time
- Byte count
- Average job rate
- Description

Windows servers must have the Agent for Windows installed on them before you can add them to the list of servers. When you add Windows servers to Backup Exec, you have the option to install the Agent for Windows to them remotely.

If you no longer want to monitor or back up a server with Backup Exec, you can remove it from the list of servers.

See [“Adding servers that you want to back up to the list of servers on the Backup and Restore tab”](#) on page 151.

See [“Removing servers from the list of servers on the Backup and Restore tab”](#) on page 152.

Adding servers that you want to back up to the list of servers on the Backup and Restore tab

Before you can create a backup definition, you must add the servers that you want to protect to the list of servers on the **Backup and Restore** tab. You can add servers during the push-installation process or you can complete the following procedure to add servers at any time.

To add servers to the list of servers

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Complete the steps to add a server or servers to the list of servers.

Note: If Backup Exec discovered servers using the **Discover Data to Back Up** option, they display on the **Browse** dialog box under the heading **Servers without an Agent for Windows installed**.

See [“Adding discovered servers to the list of servers in Backup Exec”](#) on page 715.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

Removing servers from the list of servers on the Backup and Restore tab

If you no longer want to monitor or back up a server with Backup Exec, you can remove it from the list of servers on the Backup and Restore tab. You can no longer back up or restore data from servers after you remove them from the list of servers.

Note: If you remove a server from the list and it has scheduled jobs pending, the jobs are deleted. The jobs do not run as scheduled. Do not remove a server from the list of servers if you still want to back up that server.

You should not use this procedure to delete a managed Backup Exec server from a CAS environment. You should remove managed Backup Exec servers using the **Storage** tab.

See [“Removing a managed Backup Exec server from a Backup Exec server pool”](#) on page 1366.

To remove servers from the list of servers

- 1 On the **Backup and Restore** tab, right-click the server that you want to remove from the list of servers.
- 2 Click **Remove**.
- 3 Click **Yes** to confirm that you want to remove the server from the list of servers.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

Creating a server group

Server groups are a way to organize and view server information in the list of servers. You can create server groups based on any criteria. You may want to group servers with a specific type of data or servers that reside in a specific location. Then, when you view server groups, only the server group that you select displays in the list of servers on the **Backup and Restore** tab. Viewing server groups lets you quickly monitor the status of all of the servers in the group at a glance. You can also back up an entire server group.

See [“Backing up data”](#) on page 158.

Backup Exec comes with two preconfigured server groups. The All Servers server group contains all of the servers in the list of servers. The Retired Servers server group is intended for any servers that you no longer actively monitor with Backup Exec. Servers no longer appear in the All Servers server group after you move them to the Retired Servers server group.

To create a server group

- 1 On the **Backup and Restore** tab, in the **Views** group, click **Groups**.
- 2 Click **Add**.
- 3 In the **Group name** field, type a name for the server group that you want to create. You may want to name the server group to indicate the type of servers in the group or the location at which the servers reside, for example.
- 4 In the **Description** field, type a description for the server group.
- 5 (Optional) Do any of the following to filter the servers in your environment so that you can find the servers that you want to add to the server group:

To filter servers by type	In the Server type field, select the type of server for which you want to search.
To filter servers by data type	In the Data type field, select the type of data that the server for which you want to search contains.
To filter servers by name	In the Name contains field, type all or part of the server name.

- 6 In the **Servers** group box, select the servers that you want to add to the server group, and then click **OK**.

See [“Hiding or viewing server groups on the Backup and Restore tab”](#) on page 154.

- See [“Adding servers to a server group”](#) on page 154.
- See [“Removing servers from a server group”](#) on page 155.
- See [“Editing a server group”](#) on page 155.
- See [“Moving servers to the Retired Servers server group”](#) on page 156.
- See [“Removing server groups from Backup Exec”](#) on page 157.

Hiding or viewing server groups on the Backup and Restore tab

You view server groups on the **Groups** pane. The **Groups** pane is enabled by default when you install Backup Exec. If you do not use server groups, you can hide the **Groups** pane. Double-click a server group on the **Groups** pane to view more detailed information about the server group's jobs, job history, and any active alerts.

To hide or view server groups on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Views** group, click **Groups**.
- 2 Select **Show Server Groups**.

The **Show Server Groups** option lets you hide or show the **Groups** pane to the left of the list of servers.

See [“Creating a server group”](#) on page 153.

Adding servers to a server group

You can add servers to an existing server group.

To add servers to a server group

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the group to which you want to add a server.
- 2 Select **Edit**.
- 3 In the **Servers** group box, select the servers that you want to add to the server group.

You can use the following fields to filter the list of servers so that you can find the server that you want to add:

- Server type
- Data type

- Name contains

4 Click **OK**.

See [“Creating a server group”](#) on page 153.

Removing servers from a server group

You can remove servers from an existing server group.

To remove servers from a server group

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the group from which you want to remove a server.
- 2 Select **Edit**.
- 3 In the **Servers** group box, deselect the servers that you want to remove from the server group.

You can use the following fields to filter the list of servers so that you can find the server that you want to remove:

- Server type
- Data type
- Name contains

4 Click **OK**.

See [“Creating a server group”](#) on page 153.

Editing a server group

You can edit an existing server group by changing the group's name or description.

To edit a server group

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the group that you want to edit.
- 2 Select **Edit**.

3 Do any of the following, as necessary:

To change the server group's name	In the Group name field, type the new name for the server group.
To change the server group's description	In the Description field, type the new description for the server group.

4 Click **OK**.

See [“Creating a server group”](#) on page 153.

Moving servers to the Retired Servers server group

You can retire servers from Backup Exec by moving them to the Retired Servers server group. The Retired Servers server group is intended for any servers that you no longer actively monitor with Backup Exec. You can still view any information about the retired servers on the **Backup and Restore** tab by clicking on the Retired Servers server group. However, the retired servers no longer appear in the All Servers server group with the servers that you regularly back up and monitor. It may be useful to retire servers if you use Backup Exec to monitor a large number of servers.

You cannot create new backup jobs for any servers that are in the Retired Servers server group. However, any scheduled backup jobs still run on retired servers. You can restore data from retired servers.

If you retire a server and then you decide that you want to move it back to the All Servers server group, you can click on it and drag it from the Retired Servers server group to the All Servers server group.

To retire servers from server groups

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click **Retired Servers**.
- 2 Click **Edit**.

- 3 (Optional) Do any of the following to filter the servers in your environment so that you can find the servers that you want to retire:

To filter servers by type	In the Server type field, select the type of server for which you want to search.
To filter servers by data type	In the Data type field, select the type of data that the server for which you want to search contains.
To filter servers by name	In the Name contains field, type all or part of the server name.

- 4 In the **Servers** group box, select the servers that you want to retire and then click **OK**.

See [“Creating a server group”](#) on page 153.

Moving retired servers back to the All Servers server group

You can retire servers from Backup Exec by moving them to the Retired Servers server group. The Retired Servers server group is intended for any servers that you no longer actively monitor with Backup Exec.

See [“Moving servers to the Retired Servers server group”](#) on page 156.

If you retire a server and then you decide that you want to begin monitoring it again, you can reactivate the server. You reactivate the server by moving it from the Retired Servers server group to the All Servers server group.

To reactivate retired servers

- 1 On the **Backup and Restore** tab, in the **Groups** pane, click **Retired Servers**.
- 2 In the **Servers** pane, select the server that you want to reactivate.
- 3 Drag the server to the **All Servers** group in the **Groups** pane and drop it there.

See [“Creating a server group”](#) on page 153.

Removing server groups from Backup Exec

If you no longer want to use a server group, you can remove it from Backup Exec. Removing a server group does not affect the servers in the group. It only removes the association between the servers and your ability to see them in the group.

You can still back up and monitor servers after you remove the server group to which they belong. You cannot remove either the All Servers server group or the Retired Servers server group.

To remove server groups

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the server group that you want to delete.
- 2 Click **Remove**.
- 3 Confirm that you want to remove the server group.

See [“Creating a server group”](#) on page 153.

Backing up data


When you want to back up data, you create a container that is called the backup definition.

Backup definitions contain backup selections, job templates, and stages.

Table 4-2 Backup definition contents

Item	Description
Backup selections	Backup selections include any servers, volumes, or data that you have selected to back up.
Job templates	Job templates are the collection of settings that Backup Exec uses to create jobs. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example. When you want to run a backup, Backup Exec combines the job template with the backup selections to create a backup job that runs according to the options that you specified. Backup definitions always contain one job template that uses the full backup method, but recurring jobs can also contain incremental, differential, or data-specific job templates. See “Backup methods in Backup Exec” on page 189.
Stages	Stages are optional tasks that you can run with backup jobs. Backup definitions can contain one or more stages. You can create stages that duplicate your backup data or create virtual machines with your backup data.

Figure 4-1 Backup Definition (with backup selections, full and incremental backup job templates, and a duplicate stage)



Backup Exec offers many choices for creating backup jobs to protect your data, including the following:

- Create a recurring backup definition to back up the full or partial contents of a single server or multiple servers

If you select to back up multiple servers or applications, you can create separate backup definitions for each server or application. Alternatively, you can create one backup definition that includes all servers and applications. If you create separate backup definitions, it is easier to identify problems when backup jobs fail. Also, if an issue arises with one server that causes a backup job to fail, the other backup jobs can still complete successfully. If you create one backup definition that includes all of your servers and applications, it takes less work to monitor the job status. You can select a default method for backing up multiple servers in the **Backups** dialog box in the Backup Exec Settings.

See [“Configuring the default setting for backing up multiple servers or applications”](#) on page 700.

Note: You cannot back up multiple servers with a synthetic backup job or a conversion to virtual job.

- **Create a backup definition to run only once**

A one-time backup only runs once without any recurring instances. You may want to use a one-time backup to create a baseline for a server before you upgrade it or install new software. After Backup Exec finishes running a one-time backup, it deletes the job rather than saving it with your recurring jobs. If you want to view information about a one-time backup after the job is complete, you can still view its job history.

Warning: Data lifecycle management deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup. To prevent the backup sets from being automatically deleted, you can manually retain specific backup sets or you can change the expiration date of the backup set.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

- **Create a new backup definition using an existing backup definition's settings**
If you want to create a backup definition that is similar to an existing backup definition, you can apply the existing definition's settings to a new definition. Any backup methods, job settings, and stages are copied into a new backup definition for the server or servers that you selected to back up. All that you have to do is select the backup selections. You can override any of the job settings, if necessary.
- **Create a server group out of similar computers and back up the entire group at once**
Server groups are a way to organize and view server information in the list of servers. You can create server groups based on any criteria. You may want to group servers with a specific type of data or servers that reside in a specific location.

You can also edit existing backup definitions to modify their schedules, backup selections, or other settings.

To protect remote computers, you must install the Agent for Windows on the remote computer. The Agent for Windows is a system service that runs on Windows servers and workstations. It provides efficient backup processing by

locally performing the tasks that, in typical backup technologies, require extensive network interaction.

For information about the best practices to use backup jobs, refer to the *Backup Exec Best Practices*.

See [“Methods for installing the Agent for Windows”](#) on page 68.

To back up data

- 1 Select one of the following methods to back up data:

To create recurring backup jobs

Complete the following steps:

- On the **Backup and Restore** tab, select the server, servers, or server group that you want to back up.
- Right-click the server, servers, or server group.
- On the **Backup** menu, select the backup option that you want to use.

To create one-time backups

Complete the following steps:

- On the **Backup and Restore** tab, select the server, servers, or server group that you want to back up.
- In the **Backups** group, click **One-Time Backup**.
- Select the backup option that you want to use.

To create a new backup definition using the settings from an existing backup definition

Complete the following steps:

- On the **Backup and Restore** tab, right-click the server or servers that you want to back up. You cannot reuse a backup definition's settings to create a backup job for a server group.
- On the **Backup** menu, select **Create a New Backup Using the Settings from an Existing Backup**.
- On the **Backup Job Selection** dialog box, select the backup definition that contains the settings that you want to copy.
- Click **OK**.


To create a backup definition for a server group

Complete the following steps:

- On the **Backup and Restore** tab, in the **Groups** pane, right-click the server group that you want to back up.
- On the **Backup** menu, select the backup option that you want to use.

- 2 In the **Name** field, type a name for the backup definition.

- 3 In the **Description** field, type a description for the backup definition.
- 4 In the **Selections** box, click **Edit**.



- 5 To add additional servers to the backup definition:
 - Click **Add**.
 - Select the server or servers that you want to add to the backup definition.
 - Click **OK**.

- 6 Select the data that you want to back up.

Servers are selected in their entirety by default. If you do not want to back up the entire server, double-click the server name to see all of the contents. Select the check boxes for each item that you want to back up.

Note: Deselecting a server's critical system components creates backup sets that cannot be used for some restore scenarios.

See [“About backing up critical system components”](#) on page 184.


- 7 (Optional) To prioritize the backup of specific volumes or virtual machines, select the item, and then click **Tag as business-critical**.

See [“About selecting data to back up ”](#) on page 171.

- 8 Click **OK**.

- 9 In the **Backup** box, click **Edit**.

Note: If you copied the backup settings from an existing backup definition or if you do not want to change any of the existing or default settings, you can skip to step 13



- 10 In the left pane, click **Schedule**, and then select when you want the backup job or jobs to run.
- 11 In the left pane, click **Storage**, and then select the storage device that you want to use for the backup job or jobs.

- 12 In the left pane, select any additional options that apply to the backup job or jobs.

The remaining options in the left pane are optional. The options vary depending on what you selected to back up.

For example, you might want to set up notifications so that users can receive email or texts regarding this backup definition.

Network

Select this option to specify the network interface that Backup Exec uses to access remote computers.

See [“Configuring network options for backup jobs”](#) on page 203.

Notification

Select this option to configure Backup Exec to notify specified recipients when the backup job is completed.

Each backup job and stage can be configured with different notification recipients. Backup Exec can notify people by email or text message.

See [“Notification options for jobs”](#) on page 322.

Test Run

Select this option to configure a test job that automatically tests storage capacity, credentials, and media integrity.

The test job can help you determine if there are any problems that might keep the backup job from completing successfully.

See [“Configuring automatic test run jobs for backup jobs”](#) on page 659.

Verify

Select this option to create a job that automatically verifies whether all of the data was successfully backed up when the job is completed.

A verify job can also help you determine whether the media you use is defective.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

Instant GRT

Select this option to configure the Instant GRT or the full catalog operation for any GRT-enabled jobs. You can choose to run a full catalog operation immediately after the backup job finishes, schedule the full catalog operation for another time, or run an Instant GRT operation as part of the backup job.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Advanced Open File

Select this option to configure the snapshot settings that Backup Exec uses to process the backup job. Snapshot technology lets Backup Exec capture any files that are open when a backup job runs.

You can also enable checkpoint restart, which lets you resume interrupted backup jobs.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 669.

See [“Configuring checkpoint restart”](#) on page 672.

Advanced Disk-based Backup

Select this option to configure off-host backup processing for the backup job.

See [“Configuring off-host backup options for a backup job”](#) on page 1391.

Pre/Post Commands

Select this option to configure any commands that you want to run either before the backup job begins or after the backup job is completed.

See [“Configuring pre/post commands for backup or restore jobs”](#) on page 674.

Microsoft 365

Select this option to configure parallel streams and job settings that you want to use for Microsoft 365 backups.

Files and Folders

Select this option to configure how Backup Exec processes file system attributes such as junction points and symbolic links.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

Enterprise Vault

Select this option to configure Enterprise Vault options.

See [“Enterprise Vault backup options”](#) on page 1274.

Linux and Macintosh	Select this option to configure options for any Linux or Macintosh computers that are included in the backup job. See "Linux and Unix backup options" on page 1439.
Microsoft Active Directory	Select this option to configure options for any Microsoft Active Directory data that is included in the backup job. See "Microsoft Active Directory backup job options" on page 1316.
Microsoft Exchange	Select this option to configure options for any Microsoft Exchange data that is included in the backup job. See "Backing up Exchange data" on page 1182.
Virtual Machines	Select this option to configure options for any virtual machines that are included in the backup job. See "Backing up Microsoft Hyper-V virtual machines" on page 1074. See "Backing up VMware virtual machines" on page 1013.

Performance Metrics

Select this option to collect diagnostic information and key performance indicators that provide insights about potential bottlenecks that cause performance degradation of Backup Exec operations. It provides a unified and consistent performance reporting model across the following currently supported workloads:

- VMware hosts
- Microsoft SQL Servers
- Microsoft Hyper-V hosts
- Windows files and folders

Performance metrics can be collected for backups that use disk, deduplication, and tape as storage devices.

To enable performance reporting for a backup job, select **Enable Enhanced Performance Reporting**. This option is disabled by default.

The following additional options are available:

- **Save Statistics to Resource-specific Files:**
Generates resource-specific periodic throughput statistics plotted as SVG graphs and saved to the `Backup Exec\Data` directory.
- **Sampling Interval for Periodic Statistics:** Provides precise control over the number of samples collected, depending on the job's duration. This is useful for gathering statistics from long-running jobs that generate a large volume of samples.

Microsoft SharePoint

Select this option to configure options for any Microsoft SharePoint data that is included in the backup job.

See ["Backing up Microsoft SharePoint data"](#) on page 1207.

Microsoft SQL

Select this option to configure options for any Microsoft SQL data that is included in the backup job.

See ["Backing up SQL databases and transaction logs"](#) on page 1127.

NDMP

Select this option to configure options for any NDMP data that is included in the backup job.

See ["NDMP backup options for NDMP servers"](#) on page 1403.

Oracle Select this option to configure options for any Oracle data that is included in the backup job.

See [“Oracle backup options”](#) on page 1247.

Exclusions Select this option to exclude specific files or specific types of files from backups.

See [“Excluding files from backups”](#) on page 179.

- 13 When you are finished selecting all applicable options for the backup definition, click **OK**.
- 14 (Optional) You can add stages to this backup definition to duplicate your data or create virtual machines with your backup data.

Note: You cannot add a stage to one-time backups.

See [“Adding a stage to a backup definition”](#) on page 220.

- 15 Click **OK** to save the backup definition settings.

The backup jobs and stages run at the time that you selected.

See [“Creating a server group”](#) on page 153.

See [“Editing backup definitions”](#) on page 205.

See [“About selecting data to back up ”](#) on page 171.

See [“Changing the order in which backup sources are backed up”](#) on page 176.

How Backup Exec backs up and restores volumes that are enabled with bit-locker

Backup Exec requires system or data volumes in an unlocked state during the backup and restore operation. To avoid manual changes for unlocking the bit-locked volumes, it is recommended to unlock volumes during boot or use the Auto-Unlock feature of Bit-Locker Drive Encryption. No additional configuration is required in Backup Exec.

The data backed up from the bit-locked volumes are in non-encrypted state. You can restore the data to normal volume or bit-locked volumes.

How Backup Exec backs up and restores volumes that are enabled for deduplication in Windows

Windows Server 2012 introduced native file system deduplication. When a Windows volume deduplicates data, the deduplicated data is considered to be optimized. Data in its original, non-deduplicated format is considered to be non-optimized.

Backup Exec supports the backup of volumes that have Windows deduplication enabled. If you back up optimized data, Backup Exec backs it up in its original, non-optimized format. Ensure that you have enough space to back up the non-optimized data on the backup target before you run the backup job. The optimized files are not affected on the backup source itself.

When you restore the data that you backed up, Backup Exec restores the files as non-optimized. Ensure that you have enough disk space to restore the non-optimized data. You should free space on the volume on which you want to restore the files. Run a Windows garbage collector to optimize the space on the volume before you run the restore.

How Backup Exec backs up and restores Microsoft Virtual Hard Disk (vhd) files

Microsoft Windows Hyper-V servers give users the ability to create native Virtual Hard Disk (vhd) files. A vhd file is a virtual hard disk contained in a single file. For more information about vhd files, see your Microsoft Windows documentation.

Backup Exec gives you the ability to back up and restore native vhd files. If a native vhd file is not mounted, you can back up and restore it with the volume on which it resides.

If a native vhd file is mounted to a drive letter or to an empty folder path, the file is skipped during backup jobs. You cannot include a mounted vhd as part of your backup selections. To back up the data in a mounted vhd file, select its mount point in the backup selections.

See [“Backing up data”](#) on page 158.

You can also redirect a restore job to a native vhd if you use Windows Server as a target. When you redirect a restore job to a native vhd, Backup Exec creates a vhd file that expands dynamically as you save data to it. The file expands until it reaches 2040 GB, which is the maximum size for a native vhd file. You can create one vhd file with data from all redirected backup sets or you can create a vhd file for each backup set.


Backup Exec's Agent for Hyper-V supports Microsoft vhdx files:

See “Notes about using the Agent for Hyper-V ” on page 1068.

About selecting data to back up

When you back up a server, Backup Exec includes all of the data on the server in the backup selections by default. If you want to modify the backup selections, you can click **Edit** in the **Selections** box on the **Backup Definitions Properties** dialog box.

Figure 4-2 Selections box on the Backup Definition Properties dialog box



See “Backing up data” on page 158.

Instead of backing up all of the data on a server, you can select drives, folders, files, System State, network shares, or databases on the **Browse** tab.

To expand or collapse the view for an item, click the arrow next to it or double-click the item's name. To view the contents of an item, double-click the item's icon. The item's contents appear in the right frame of the backup selections view. You can traverse file levels from either side of the window by clicking folders and subfolders as they appear.

Windows server provides the ability to encrypt volumes using the Bit locker functionality. Backup Exec supports the backup and restore of Bitlocker enabled volumes either by manually unlocking the destination volumes or by using the Microsoft auto-unlock functionality. For more information, refer to the Microsoft documentation.

When you browse remote selections, Backup Exec requires a valid logon account to expand the computer contents. If the default logon account does not enable access to a remote selection, Backup Exec prompts you to select another existing logon account. You can also create a new logon account that can access the selection.

See [“Testing credentials for backup sources”](#) on page 212.

See [“Replacing the credentials for a backup source”](#) on page 213.

See [“Creating new credentials for a backup source”](#) on page 214.

To include data in the backup, select the check box next to the drive or directory that you want to back up.

This topic includes the following information:

[Tagging resources as business-critical](#)

[Including and excluding specific files or subdirectories](#)

[Selecting multiple servers or applications](#)

[Selecting critical system components](#)

[Using fully qualified domain names in backup selections](#)

Tagging resources as business-critical

The **Tag as Business-Critical** option lets you select which resources are most important to you. When an item is tagged as a business-critical resource, Backup Exec prioritizes the backup of that resource within the backup window before any resources that are not tagged as business-critical. If you tag a resource as business-critical, all of its children are also tagged as critical. When you tag an item as a business-critical resource, Backup Exec changes the icon for that resource in the backup set view and in the restore wizard.

The **Tag as Business-Critical** option is enabled for the following resources:

- Windows volumes and UNIX root volumes
- Exchange databases
- SQL instances
- Virtual machine folders
- Virtual machines
- SharePoint resources
- Oracle instances

Note that the **Tag as business-critical** option is disabled on partial selections. In addition, the option is not applicable to Enterprise Vault resources, such as Compliance Accelerator or Discovery Accelerator.

When an item is tagged as a business-critical resource, it applies only to that job. It is not a global setting. When you tag a server's resource as business-critical and also select other resources from the same server, then the backup priority for the business-critical resource is evaluated in relation to the other resources for that server only. For multi-server jobs, business-critical resources are backed up by the order in which the servers are prioritized, with the business-critical resources for servers at the top of the list being backed up before business-critical resources for the servers that are lower on the list.

Critical system resources can be tagged as business-critical as well. If critical system resources are not selected, then the resources that are tagged as business-critical are backed up before critical system resources if they are in the same backup job. However, the best practice for critical system resources is to back them up in a separate backup job. If you include both business-critical resources and critical system resources in the same backup job, note that system state is always backed up last regardless of whether it was tagged as a business-critical resource or not.

After you tag items as business-critical resources you can use the arrow buttons on the **Selection Details** tab to change the order in which the items are backed up. However, an item that is not tagged as a business-critical resource cannot be moved ahead of an item that is tagged as a business-critical resource. For example, if you tagged your C drive as a business-critical resource, but did not tag your E drive as a business-critical resource, you cannot move the E drive ahead of the C drive.

See [“Changing the order in which backup sources are backed up”](#) on page 176.

Note: During a rolling upgrade, the tag as business-critical feature is not applied to jobs that are delegated to managed Backup Exec servers until the managed servers are upgraded to the most recent version of Backup Exec.

Including and excluding specific files or subdirectories

The **Selection Details** tab lets you include or exclude files for backups by specifying file attributes. Exclusions apply to all of the jobs in a backup definition.

You can do any of the following on the **Selection Details** tab:

- Include or exclude subdirectories. For example, you can choose to back up a parent folder without backing up any folders that reside inside it.
- Include only modified files. For example, you can choose to back up only the files that have changed since the last backup job.
- Include only read-only files.
- Include or exclude files by file name attributes. For example, you can select only files with .txt extensions, or exclude files with .exe extensions from a backup. If you exclude files by an attribute that does not exist, all files of that type are excluded. For example, excludes based on SQL database dates result in global SQL excludes since SQL databases do not have date attributes.
- Select only any files that fall within a specified date range. For example, you can select any files that were created or modified during the month of December.
- Specify the files that have not been accessed in a specified number of days. For example, you can select the files that have not been accessed in 30 days from your "My Documents" folder. Then, run a full backup job for which you select the method to back up and delete the files.

See [“Excluding files from backups”](#) on page 179.

See [“Including specific files with a backup job's backup selections”](#) on page 182.

Selecting multiple servers or applications


If you select to back up data from multiple servers or applications, you can create separate backup definitions for each server or application. Alternatively, you can create one backup definition that includes all servers and applications. If you create separate backup definitions, it is easier to identify problems when backup jobs fail. Also, if an issue arises with one server that causes a backup job to fail, the other backup jobs can still complete successfully. If you create one backup definition that includes all of your servers and applications, it takes less work to monitor the job's status. You can select a default method for backing up multiple servers in the **Backups** dialog of the Backup Exec Settings.

See [“Configuring the default setting for backing up multiple servers or applications”](#) on page 700.

Selecting critical system components

When all the critical system components are included in your backup job selections, the **Simplified Disaster Recovery** indicator on the selections pane reads **ON**. If you deselect one or more critical system component files, the indicator changes to **OFF**.

Figure 4-3 Simplified Disaster Recovery indicator is ON in the Backup Definition Properties dialog box



See [“About backing up critical system components”](#) on page 184.

If you deselect any critical system components, it can disqualify your backup data from being used in certain types of restore scenarios.

You must include all critical system components in your backup selections if you intend to use any of the following restore scenarios:

- Simplified Disaster Recovery
- Conversion to virtual machines
- Complete online restore of a Microsoft Windows computer

Using fully qualified domain names in backup selections

You can enter fully qualified domain names in Backup Exec anywhere that you can enter a computer name. In addition, Backup Exec can show fully qualified domain names where computer names are listed.

For fully qualified domain names, the following rules apply:

- The maximum number of characters for each label (the text between the dots) is 63
- The maximum total number of characters in the fully qualified name is 254, including the dots, but excluding the \\
- The name cannot include the following characters: * | < > ?

It is not recommended using both fully qualified domain names and non-qualified domain names. It is recommended that you use fully qualified domain names.

For example, if you have a computer named Test_Computer, you can have two selections for it. One selection is called Test_Computer. The fully qualified selection is called Test_Computer.domain.company.com. In this case, Backup Exec treats each selection as a separate computer, even though both selections are for the same computer. For any backup jobs that use the short computer name, the catalog contains the short computer name. For any backup jobs that use the fully qualified name, the catalog contains the fully qualified name.

See [“Changing the order in which backup sources are backed up”](#) on page 176.

Changing the order in which backup sources are backed up

After you make selections for a backup job, you can configure Backup Exec to process those selections in a certain order. You can create backup jobs where your most important backup sources are backed up first, for example.

You must select a backup source to reorder the source or any of its children. If a backup source is included dynamically in the backup, you cannot specify the

order in which any of its children are backed up. Any time that you include or exclude a backup source, it creates an entry in the backup's selection details. You can reorder any of the entries that appear on the **Selection Details** tab.

There are some limitations to which backup sources can be reordered:

- Any backup sources that reside on branches in the browse tree must be backed up together. Branches are containers for backup sources in the browse tree. Individual applications appear in the Backup Exec browse tree as branches. If you have a server which contains SharePoint, SQL, and Exchange data, each application appears as a branch. You can change the backup order of individual items on a branch, but you cannot change the order of items across multiple branches.

For example, if you want to back up a volume and a SQL database, each of the SQL instances must be backed up in succession. You can change the order in which the SQL instances are backed up. You cannot back up a SQL instance and then the volume before backing up the other SQL instances, however.


- All of the data that resides on a specific backup source is treated as a single item. You can change the order in which the backup sources are backed up in relation to one another. You can also change the order in which the data on a backup source is backed up. However, you cannot reorder the data across multiple backup sources.

For example, if you back up two volumes, you can select which volume should be backed up first. You can also select the order in which the data on each volume is backed up. However, you cannot back up some of the data from one volume and then the other volume before backing up the rest of the data on the first volume.

- Some system resources cannot be reordered. For example, Shadow Copy Components and System State backup sources must always be backed up last.

To change the order in which backup sources are backed up

- 1 Create or edit a backup definition.
- 2 In the **Selections** box, click **Edit**.



- 3 Do one of the following:

To reorder the servers in a backup definition with multiple servers

In the left pane of the **Backup Selections** dialog box, use the up and down arrows to configure the order of the servers or applications.

Backup Exec backs up the servers or applications at the top first.

To reorder the backup sources within a server

Do the following:

- In the left pane of the **Backup Selections** dialog box, select the server that contains the resources that you want to reorder.
- In the right pane of the **Backup Selections** dialog box, select the **Selection Details** tab.
- To reorder the backup sources when a server is selected in its entirety, you must insert the individual selections that you want to reorder:
 - Click **Insert**, and then select **Insert Selection**.
 - Select the backup source that you want to reorder, and then click **OK**.

Repeat this step for each backup source that you want to reorder.

- In the right pane of the **Backup Selections** dialog box, use the up and down arrows to configure the order of the backup sources.

Backup Exec backs up the backup sources at the top first.

- Click **OK**.

4 Click **OK**.

See [“Backing up data”](#) on page 158.

See [“About selecting data to back up ”](#) on page 171.

Excluding files from backups

You can exclude specific files or types of files from backups if you want to ensure that they do not get backed up. For example, you may not want to back up any mp3 files, read-only files, or files from specific directories.

You can exclude files from backups in any of three ways:

- Exclude files from backup jobs

When you exclude files from individual backup jobs, the exclusions do not affect any of the other backup jobs in the related backup definition. The

job-level exclusions also do not affect any of your other backup definitions. The exclusion applies only to the backup job to which you apply it. You cannot apply a job-level exclusion to Full backups.

- **Exclude files from backup definitions**
When you exclude files from backup definitions, the exclusions do not affect any of your other backup definitions. The exclusions apply only to any jobs in that backup definition.
- **Exclude files globally from all backups**
When you globally exclude files from all backups, the exclusions apply to all of your backup definitions. Backup Exec automatically skips any global exclusions when you run backup jobs.

To exclude files from backups

- 1 Do any of the following:

To exclude files from a backup job

Complete the following steps:

- Create or edit a backup definition.
- In the **Backup** box, click **Edit**.
- In the left pane, select **Exclusions**.
- Click **Insert**.

To exclude files from a backup definition

Complete the following steps:

- Create or edit a backup definition.
- In the **Selections** box, click **Edit**.
- On the **Selection Details** tab, click **Insert** and then select **Add Backup-Level Exclusion**.

- To globally exclude files from all backups

Complete the following steps:

 - Click the Backup Exec button.
 - Select **Configuration and Settings**, and then select **Job Defaults**.
 - Click **Exclude Selections**.
 - Click **Insert**.
- 2

Complete any of the following fields to identify the files that you want to exclude:

Resource name	Enter the name of the volume or drive that you want to exclude from backups.
Path	Enter the path of the folder and/or subfolder that contains the files that you want to exclude. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.
Name	Enter the name of a specific file that you want to exclude from the backup. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters. For example, to exclude all files with a .exe extension, type "***.exe".

3 Complete any of the following options, as necessary:

Apply to subdirectories	Select this option to exclude the contents of all of the subfolders when you select a directory.
Only modified files	Select this option to exclude only the files that have been modified in the directory that you specified.
Only read-only files	Select this option to exclude only read-only files in the directory that you specified.
Files dated	Select this option to exclude the files that were created or modified during a specific time period. You must select the beginning and ending dates for the time period.
Files not accessed in X days	Select this option to exclude any files that have not been accessed in a specific number of days. You must enter the number of days.

4 Click **OK**.

See [“About selecting data to back up ”](#) on page 171.

See [“Including specific files with a backup job's backup selections”](#) on page 182.

Including specific files with a backup job's backup selections

When you create a backup definition, you create a list of backup selections. The backup selections are the files and data that you want to back up when the backup jobs run. You can include additional specific files or backup sources with the rest of your backup selections. When you include files with a job's backup selections, you can select certain conditions which must be met for the file to be included. For example, you may want to include any read only files that reside in a specific directory. When you include a file with a backup job's backup selections, Backup Exec backs up the regular backup selections and it includes the file or files that you specifically selected.

To include specific files with a backup job's backup selections

- 1 Create or edit a backup definition.
- 2 In the **Selections** box, click **Edit**.
- 3 On the **Selection Details** tab, click **Insert** and then select **Insert Selection**.

- 4 Complete any of the following fields to identify the files that you want to include:

Resource name	Enter the name of the volume or drive that you want to include in the backup.
Path	Enter the path of the folder and/or subfolder that contains the files that you want to include. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.
Name	Enter the name of a specific file that you want to include in the backup. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters. For example, to include all files with a .exe extension, type "***.exe".

- 5 Complete any of the following options as necessary:

Apply to subdirectories	Select this option to include the contents of all of the subfolders when you select a directory.
Only modified files	Select this option to include only the files that have been modified in the directory that you specified.
Only read-only files	Select this option to include only read-only files in the directory that you specified.
Files dated	Select this option to include the files that were created or modified during a specific time period. You must select the beginning and ending dates for the time period.
Files not accessed in X days	Select this option to include any files that have not been accessed in a specific number of days. You must enter the number of days.
Tag as business-critical	Select this option to prioritize the backup of the selected resource above other resources that are not tagged as business-critical.

- 6 Click **OK**.

See [“About selecting data to back up”](#) on page 171.


See [“Excluding files from backups”](#) on page 179.

About backing up critical system components

Backup Exec is configured to automatically back up the critical system components that you need to perform a full system restore. Backing up critical system components ensures that you are capable of recovering your computers in the event of a disaster.

When all the critical system components are included in your backup job selections, the **Simplified Disaster Recovery** indicator on the selections pane reads **ON**. If you deselect one or more critical system component files, the indicator changes to **OFF**.

Figure 4-4 Simplified Disaster Recovery indicator is ON in the Backup Definition Properties dialog box



You must include all critical system components in your backup selections if you intend to use any of the following restore scenarios:

- Simplified Disaster Recovery
See [“About Simplified Disaster Recovery”](#) on page 877.
- Conversion to virtual machines
See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 589.
- Complete online restore of a Microsoft Windows computer

See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 246.

Some restore scenarios are only available for certain data types and Backup Exec agents. Additionally, you must select a specific backup method for some data types if you intend to use a physical-to-virtual conversion or a backup-to-virtual conversion.

Table 4-3 Restore scenarios by data type

File System or Agent Name	Operating System and Applications Are Installed and Functional	Complete Online Restore Can Be Performed	Physical-to-Virtual Conversion Can Be Performed	Backup-to-Virtual Conversion Can Be Performed	Simplified Disaster Recovery Can Be Performed
Active Directory	Yes	Yes	Yes	Yes	Yes
Active Directory Lightweight	Yes	No	No	No	No
CSV	Yes	Yes	No	No	Yes
EFI	Yes	Yes	No	No	Yes
Enterprise Vault	Yes	No	No	No	No
Exchange Agent	Yes	No	Yes, any incremental or differential backups must use the block-level backup method	Yes, any incremental or differential backups must use the block-level backup method	No
FAT	Yes	Yes	Yes	Yes	Yes
Hyper-V Agent	Yes	No	No	No	No
NTFS	Yes	Yes	Yes	Yes	Yes

Table 4-3 Restore scenarios by data type (*continued*)

File System or Agent Name	Operating System and Applications Are Installed and Functional	Complete Online Restore Can Be Performed	Physical-to-Virtual Conversion Can Be Performed	Backup-to-Virtual Conversion Can Be Performed	Simplified Disaster Recovery Can Be Performed
Oracle RMAN Windows Agent	Yes	No	No	No	No
ReFS	Yes	Yes	No	No	Yes
Shadow Copy Components	Yes	Yes	Yes	Yes	Yes
SharePoint Agent	Yes	No	No	No	No
SQL Agent	Yes	No	Yes, any incremental or differential backups must use the block-level backup method	Yes, any incremental or differential backups must use the block-level backup method	No
System State	Yes	Yes	Yes	Yes	Yes
Utility Partition	Yes	Yes	No	No	Yes
VMware Agent	Yes	No	No	No	No
Windows Deduplication	Yes	Yes	No	No	Yes

You cannot individually select critical system components as backup selections. You must select the entire server to ensure that all critical system components are included in the backup. When you select to back up a server, Backup Exec includes all of the server's system devices and application agents. Backup Exec

dynamically discovers and protects all critical and non-critical system devices and application agents.

You can explicitly exclude any non-critical devices or application data from the backup selections without affecting your ability to perform a full system restore. You can exclude Microsoft Exchange data from your backup, for example, and still use the backup sets to perform a disaster recovery.

The following system resources are considered critical and they must be included in backups if you want to be able to use the backup sets to perform a full system restore:

- System volume (including EFI and utility partitions)
- Boot volume (excluding operating system)
- Services application volumes (boot, system, and automatic startup)
- System State devices and volumes (including Active Directory, System Files, etc.)
- Windows Recovery Partition (WinRE) on any applicable versions of Windows

For information about the best practices to use Backup Exec to protect critical system components, refer to *Backup Exec Best Practices*.

Backup Exec Shadow Copy Components file system

The Backup Exec Shadow Copy Components file system uses Microsoft's Volume Shadow Copy Service to protect critical operating system and application service data, and third-party application and user data on Windows Servers.

Volume Shadow Copy Service allows a computer to be backed up while applications and services are running by providing a copy of a volume when a backup is initiated. Applications do not need to be shut down to ensure a successful volume backup. Volume Shadow Copy Service enables third-party vendors to create snapshot plug-ins, or Writers, for use with this shadow copy technology.

A Writer is specific code within an application that participates in the Volume Shadow Copy Service framework to provide point-in-time, recovery-consistent operating system and application data. Writers appear as Shadow Copy Components, which are listed as data in backup and restore selections.

Only Writers that have been tested for use with Backup Exec are available for selection in the backup selections. Other Writers may be displayed in the selections, but they cannot be selected for backup.

If you select a volume that contains Shadow Copy data for backup, Backup Exec determines which Shadow Copy files should not be included in a volume-level

backup. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as "in use - skipped." If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

The Windows SharePoint Services feature pack utilizes a SQL (MSDE) instance called SHAREPOINT as a repository for shared information and collaboration data. In the absence of a SQL Agent installation, the SQL SHAREPOINT instance can be protected by the Shadow Copy Components file system. If the SQL Agent is installed, then the SQL SHAREPOINT instance can be protected by the SQL Agent.

Note: If Windows SharePoint Services is installed using an instance name other than the default SHAREPOINT instance name, then it cannot be protected by the Shadow Copy Components file system. In that case, the SQL Agent must be used to protect the SQL SHAREPOINT instance.

Backup methods in Backup Exec

When you create a backup definition, you must select backup methods for each backup job. The backup method determines what data Backup Exec backs up. The standard backup methods are full, differential, and incremental. However, you can also select specific types of backup methods for some agents and types of data.

Each backup definition must contain one backup job that uses the full backup method. This initial full backup job establishes a baseline by backing up all of the data that you selected. Then, you can add additional differential or incremental backup jobs to the backup definition. You can select to use data-specific backup methods instead of the standard backup methods by selecting the appropriate backup method on the options page for that agent.

This topic includes the following information:

[Full backup method](#)

[Differential backup method](#)

[Incremental backup method](#)

[Backup method advantages and disadvantages](#)

[Backup methods for specific types of data](#)

Full backup method

Full backups include all of the data that was selected for backup. Backup Exec detects that the server was backed up.

Note: You should perform a full backup of your server to establish a baseline for disaster recovery.

Duplicate backups, which include all selected data, are a type of full backup. Duplicate backups do not affect any tape media rotation strategy because the archive bit is not reset.

It is recommended that you always run full backups before and after you upgrade, update, or modify any backup sources. This recommendation applies to any significant operating system and application configuration or modifications.

Table 4-4 Backup sources and scenarios for which a full backup is recommended

Backup source	Scenarios for which a full backup is recommended
Operating system	You should run a full backup when you: ■ Upgrade to a new version of the operating system. ■ Update the existing operating system using Windows Update. Note: You should run a full backup before you apply the update. Apply the update and then restart the computer. Then you should run another full backup. ■ Add, modify, or remove any server roles or features.
Applications	You should run a full backup when you: ■ Install new applications. ■ Upgrade applications to new versions. ■ Update the existing version of applications.

Table 4-4 Backup sources and scenarios for which a full backup is recommended (*continued*)

Backup source	Scenarios for which a full backup is recommended
Backup Exec	You should run a full backup when you: ■ Upgrade to a new version of Backup Exec. Note: Any recurring jobs that are migrated to a new version of Backup Exec retain their existing schedules. You should manually run a full backup for any existing jobs before any incremental backups or differential backups run. ■ Update the existing version of Backup Exec using Arctera Update. ■ Make configuration changes.

Running full backups before and after each of these scenarios helps to ensure that you can restore back to the previous configuration, if necessary.

Differential backup method

Differential backups include all files that have changed since the last full backup or incremental backup. The difference between differential and incremental backups is that differential backups are cumulative. After a differential backup, each subsequent differential backup backs up the same files as the previous differential backup. It also backs up as any new files or changed files dating back to the last full backup or incremental backup.

Note: In a backup definition that includes a differential task, all of the backup tasks must use storage devices that the same Backup Exec server can access.

By default, Backup Exec uses the Windows Change Journal to determine if files were previously backed up. You can also configure Backup Exec to use a file's modified time or archive bit to determine if the file was backed up.

Note: You cannot use Backup Exec catalogs to determine if files were backed up for any differential backups.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

Differential backups allow much easier restoration of an entire device than incremental backups since fewer backups are required. Using fewer media also decreases the risk of having a restore job fail because of media errors.

Incremental backup method

Incremental backups include only the files that have changed since the last full or incremental backup. The difference between incremental and differential backups is that incremental backups are not cumulative. Each incremental backup creates a baseline. After the incremental backup, the subsequent incremental backup or differential backup backs up only any new files or changed files dating back to the baseline.

Note: In a backup definition that includes an incremental task, all of the backup tasks must use storage devices that the same Backup Exec server can access.

By default, Backup Exec uses the Windows Change Journal to determine if files were previously backed up. You can also configure Backup Exec to use a file's modified time, archive bit, or the Backup Exec catalogs to determine if the file was backed up.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

Incremental backups take much less time than full or differential backups to complete. They also require less storage space for backed up data since only any files that have changed since the last backup are backed up.

Backup method advantages and disadvantages

Each backup method has advantages and disadvantages.

Table 4-5 Backup method Advantages and Disadvantages

Method	Advantages	Disadvantages
Full	■ Files are easy to find Full backups include all the data that you selected to back up. Therefore, you don't have to search through several backup sets to find a file that you need to restore. ■ A current backup of your entire system is available on one backup set If you run a full backup of your entire system and then need to restore it, all of the most current information is located in one place.	■ Redundant backups Most of the files on your file server do not change. Each full backup that follows the first is merely a copy of what has already been backed up. Full backups require more storage. ■ Full backups take longer to perform Full backups can be time consuming, especially when you have other servers on the network that need to be backed up (for example, agent workstations, remote servers).

Table 4-5 Backup method Advantages and Disadvantages (*continued*)

Method	Advantages	Disadvantages
Differential	- Files are easy to find Restoring a system that is backed up with a differential method requires fewer backups. Differentials require the latest full backup, any subsequent incremental backups, and the latest differential backup. Restoring differentials is less time consuming than restoring incrementals. Restoring incrementals requires the latest full backup and all incremental backups that were created since the full backup. - Less time is required for backup and restore Differential backups take less time to restore than full backups. Faster recovery is possible in disaster situations because you only need backup sets from the latest full backup, any subsequent incremental backups, and the latest differential backup to fully restore a server.	Redundant backups All of the files that were created or modified since the last full backup are included; thus creating redundant backups.
Incremental	- Better use of storage Only the files that have changed since the last backup are included, so much less data storage space is required. - Less time is required for backup Incremental backups take much less time than full and differential backups to complete.	Backups are spread across multiple backup sets Since multiple backup sets are required in a disaster situation, recovering a server can take longer. In addition, the backup sets must be restored in the correct order to effectively bring the system up to date.

Consider the following backup strategy scenario:

You want to implement a backup strategy for the office file server. All backup strategies begin with a full backup (the backup of an entire server using the full backup method). So you create and submit a full backup job to run at the end of the day on Friday.

Most files on the server, such as operating system files and application files seldom change. Therefore, you decide that you can save time and storage by using incremental backups or differential backups. You opt to use incremental backups. You schedule a job to run at the end of each day, Monday through Thursday, with the incremental backup method.

On Friday, your backup sets contain all of the data on the file server. Backup Exec changes all of the files' statuses to backed up. At the end of the day on Monday, the incremental job runs and only the files that were created or changed are backed up. When the incremental job completes, Backup Exec turns off the archive bit, showing that the files have been backed up. On Tuesday through Thursday, the same events happen.

If your file server then crashed on Thursday, you would restore each backup in the order in which it was created. You would begin with Friday's backup and proceed through Wednesday's backup.

If you had decided to perform differential backups on Monday through Thursday, you would have only needed Friday's and Wednesday's backup sets. Friday's backup sets would have included all of the data from the original backup. Wednesday's backup sets would have included every file that had been created or changed since Friday's backup.

Backup methods for specific types of data

Agents and features may have specific types of backup methods.

By default, each backup definition contains a backup job that limits the available backup method to the initial full backup for most data types. You can configure additional backup jobs that contain special backup methods for specific types of data. You can name the job templates for those backup jobs so that they have additional meaning to your organization's overall backup strategy.

See [“Configuring backup methods for backup jobs”](#) on page 197.

When you create a backup definition that contains more than one backup method for multiple data types, it is called a mixed backup in the Job Monitor.

Table 4-6 Available Backup Methods By Data Type

Data type	Job type and backup method
Files and Folders	Initial full: ■ Full - Back up files Additional backup methods for files and folders: ■ Full - Back up files ■ Full Copy - Back up files (copy) ■ Differential - Back up changed files since the last full ■ Incremental - Back up changed files since the last full or incremental
Enterprise Vault	Initial full: ■ Full - Back up components Additional backup methods for Enterprise Vault: ■ Full- Back up components ■ Differential - Back up component changes since the last full ■ Incremental - Back up component changes since the last full or incremental
Microsoft Exchange	Initial full: ■ Full - Back up databases and logs (truncate logs) ■ Full Copy - Back up databases and logs Additional backup methods for Microsoft Exchange: ■ Full - Back up databases and logs (truncate logs) ■ Full Copy - Back up databases and logs ■ Differential - Back up logs ■ Incremental - Back up logs (truncate logs)
Virtual Machines	Initial full: ■ Full - Back up virtual machines Additional backup methods for virtual machines: ■ Full - Back up virtual machines ■ Differential - Back up virtual machine changes since the last full ■ Incremental - Back up virtual machine changes since the last full or incremental

Table 4-6 Available Backup Methods By Data Type (*continued*)

Data type	Job type and backup method
Microsoft SharePoint	Initial full: ■ Full - Back up databases ■ Full Copy - Back up databases (copy) Additional backup methods for Microsoft SharePoint: ■ Full - Back up databases ■ Full Copy - Back up databases (copy) ■ Differential - Back up database changes since the last full ■ Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job ■ Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job ■ Log - Back up and truncate transaction log
Microsoft SQL	Initial full: ■ Full - Back up databases ■ Full Copy - Back up databases (copy) Additional backup methods for Microsoft SQL: ■ Full - Back up databases ■ Full Copy - Back up databases (copy) ■ Automatic - Back up transaction log if enabled and then back up database changes since the last full or incremental ■ Log - Back up and truncate transaction log ■ Log No Truncate - Back up without truncating transaction log ■ Differential - Back up database changes since the last full ■ Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job ■ Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job ■ Database Snapshot - Read-only point-in-time copy of databases

Table 4-6 Available Backup Methods By Data Type (*continued*)

Data type	Job type and backup method
NDMP (all)	Initial full: ■ Level 0 - Full backup Additional backup methods for NDMP: ■ Level 0 - Full backup ■ Level 1 - Incremental (backs up new or modified files since level 0) ■ Level 2 - Incremental (backs up new or modified files since level 1) ■ Level 3 - Incremental (backs up new or modified files since level 2) ■ Level 4 - Incremental (backs up new or modified files since level 3) ■ Level 5 - Incremental (backs up new or modified files since level 4) ■ Level 6 - Incremental (backs up new or modified files since level 5) ■ Level 7 - Incremental (backs up new or modified files since level 6) ■ Level 8 - Incremental (backs up new or modified files since level 7) ■ Level 9 - Incremental (backs up new or modified files since level 8)
Oracle	Initial full: ■ Full - Back up selections Additional backup methods for Oracle: ■ Full - Back up selections ■ Differential - Back up changes since the last full ■ Incremental - Back up changes since the last full or incremental

Configuring backup methods for backup jobs

When you create a backup definition, you must select backup methods for each backup job. The backup method determines what data Backup Exec backs up. The standard backup methods are full, differential, and incremental. However, you can also select specific backup methods for some agents and types of data.

See [“Backup methods in Backup Exec”](#) on page 189.

By default, each backup definition contains one backup job with an initial full backup method that cannot be changed and a backup job with the incremental backup method. You can replace the incremental job with a differential job, add additional incremental or differential backup jobs, or change the backup methods to data-specific backup methods for certain types of data. Each backup method can be given a unique name to help make it more easily identifiable.

To configure backup methods for backup jobs

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 (Optional) To add additional backup jobs to the backup definition, complete the following steps:
 - Click **Add a Backup Job**.
 - Select the type of backup method that you want to use for the new backup job.

Note: You can change the backup method to a data specific one for specific agents and types of data later.

- 5 Type a name for each backup job that you want to configure a backup method for in the **Job template name** field.
- 6 (Optional) To change any of the backup methods that you selected into a data-specific backup method, complete the following steps:
 - In the left pane, select the agent or data type for which you want to configure the data-specific backup method.
For example, if you want to select a data-specific backup method for SQL data, select **Microsoft SQL**.
 - In the **Backup method** field, select the data-specific backup method for each applicable backup job.
- 7 Click **OK**.

How Backup Exec determines if a file has been backed up

If you use the incremental or the differential backup method as part of your backup strategy, Backup Exec must know when a file has been modified. Full

backups include all of the data that you selected to back up. Subsequent incremental and differential backups back up only new files and any files that have changed.

Whenever a file is created or changed, a computer's file system notes and records the change. You can specify the method by which Backup Exec determines if a file needs to be backed up in the **Files and Folders** options when you create a backup job. Backup Exec uses the method that you choose to determine when a file is created or changed.

This topic includes the following information:

[Modified time](#)

[Archive bit](#)

[Catalogs](#)

[How Backup Exec uses the change journal to determine which files have changed](#)

[Resolving change journal errors](#)

Modified time

When Backup Exec runs a full backup or incremental backup job, the time that the backup job starts is recorded in the Backup Exec Database. Backup Exec adds the time of the backup job to the Backup Exec Database only if the full backup job completes successfully. The next time that you run an incremental backup job or a differential backup job, Backup Exec compares the file system time to the backup time. If the file system time is later than the time that is recorded in the database, the file is backed up. If the file's modified time is older than the previous backup's modified time, that file is not backed up. If the job does not complete successfully, subsequent differential or incremental backup jobs back up all of the data instead of only the data that has changed.

Note: A file's last modified date and timestamp do not change when the file is copied or moved. To ensure that the files are protected, run a full backup after you copy or move files.

When you run an incremental backup job, Backup Exec records a new time in the Backup Exec Database. The database time is not updated for differential backup jobs.

When you select the modified time method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to the previous backup time to determine if the file has changed.

Using modified time lets Backup Exec run more accurate incremental backups or differential backups even if other processes have modified files' archive bits.

Archive bit

Backup Exec uses the archive bit from the file system to determine if a file has changed since the last time it was backed up.

When you use the archive bit, Backup Exec turns the archive bit off when a file is backed up. Turning off the archive bit indicates to Backup Exec that the file has been backed up. If the file changes again before the next backup job, the bit is turned on again. Backup Exec backs up the file in the next backup.

If the next backup job is a full backup job or an incremental backup job, the bit is turned off when the backup job completes. If the next backup job is a differential backup job, the archive bit is left intact.

Catalogs

Backup Exec compares path names, modified time, deleted and renamed files and folders, and other attributes. When you select the catalog method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to previous catalogs to determine if it has changed.

The catalog method is only available if the Advanced Disk-based Backup feature (ADBO) is installed.

Note: The off-host backup feature of ADBO does not support the catalog method.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

How Backup Exec uses the change journal to determine which files have changed

When you choose the modified time backup method or the catalog backup method for files, Backup Exec uses the update sequence number (USN) change journal. Backup Exec scans the change journal to get a list of the changed files to back up, rather than scanning all files. Use of the change journal reduces the amount of time that Backup Exec requires to perform an incremental backup or a differential backup.

You cannot disable the use of the change journal for either the modified time backup method or the catalog backup method. If the change journal is not already enabled on the volume, then Backup Exec automatically enables it. The volume must support change journals. For example, NTFS and ReFS volumes support change journals, but FAT volumes do not.

When Backup Exec enables the change journal, it sets the change journal defaults as shown in [Table 4-7](#). Backup Exec does not modify settings for a change journal that already exists.

Table 4-7 Change journal defaults set by Backup Exec

Volume size	Defaults
128 GB or larger	■ Change Journal Max Size: 32 MB ■ Allocation Delta: 4 MB
Over 64 GB - 127 GB	■ Change Journal Max Size: 16 MB ■ Allocation Delta: 2 MB
Over 4 GB - 63 GB	■ Change Journal Max Size: 8 MB ■ Allocation Delta: 1 MB
4 GB or less	■ Change Journal Max Size: 1 MB ■ Allocation Delta: 256 KB

Note: The Windows Boot Volume is preconfigured by Windows with default values.

Resolving change journal errors

If change journal errors occur, then Backup Exec uses the modified time backup method or the catalog backup method without using the change journal.

When change journal errors occur, you can try the following solutions:

- Select the **Using archive bit** backup method for files.
- Select the file backup method that was not in use when the change journal errors occurred. For example, if you selected the **Modified time** method and change journal errors occurred, then select the **Using catalogs** method and run the backup again.
- Increase the size of the change journal database.
Refer to your Microsoft documentation for information on how to use fsutil to query, create, modify, or delete the change journal.
- Avoid using the local Backup Exec server to back up itself when the Backup Exec server is installed on the Windows Boot volume. Instead, use a remote Backup Exec server to back up the local server, or install the Backup Exec server on a data volume.

Note: The off-host backup feature of ADBO does not support the catalog backup method for files.

Configuring Backup Exec to automatically delete files after a backup

When you run a full backup, you can select to back up the files and then delete them. The **Delete selected files and folders after successful backup** option lets you free disk space on your server by deleting files and folders from the server after they are successfully backed up. Backup Exec backs up the selected data, verifies the backup sets, and then deletes the data from the server. You can back up and delete files for full backups only.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

Backup Exec performs a verify operation after the data is backed up. If the verify operation fails, the job stops and you are notified. If you get a verification failure, view the job log. Try to correct the problem, and then retry the job. After the data is backed up and verified, Backup Exec deletes the selected data. The job log contains a list of the data that is deleted.

You can enable the checkpoint restart option for a full backup job that uses the **Delete selected files and folders after successful backup** option. If the job fails and is resumed, the files are not deleted from the source volume after the backup completes.

To configure Backup Exec to automatically delete files after a backup

1 Do one of the following:

To configure Backup Exec to automatically delete files after all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure Backup Exec to automatically delete files.

To configure Backup Exec to automatically delete files for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Files and Folders**.

- 3 Select **Delete selected files and folders after successful backup**.

Note: You can select to delete files and folders only for full backups.

- 4 Click **OK**.

Configuring network options for backup jobs

You can configure options for how Backup Exec works with your network. Backup Exec contains global network and security settings that apply to all jobs.

See [“Changing network and security options for Backup Exec”](#) on page 718.

You can override the global network settings when you create backup jobs if the global settings do not apply in a particular instance. Complete the steps in the following procedure to configure network options for individual backup jobs.

Note: Network options differ in CAS environments.

To configure network options for backup jobs

- 1 Create a new backup definition or edit an existing backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Network**.
- 4 Complete the following options:

Network interface	Select the name of the network interface card that connects the Backup Exec server to the network that you want to use for this backup job. The list includes all available network interfaces on the Backup Exec server.
Protocol	Select the protocol you want to use for this backup job. The options are as follows: ■ Use any available protocol ■ IPv4 ■ IPv6
Subnet	Select the 32-bit number that determines the subnet to which the network interface card belongs.

Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	Select this option to let Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network. If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.
Interface Details	Click this option to view the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.
Allow managed Backup Exec server to use any network interface to access Backup Exec agents	Select this option to let a job use any network interface to access Backup Exec agents if the selected network interface is unavailable. Enabling this option lets the managed Backup Exec server use an alternate network interface to run any important backup jobs that would otherwise fail. This option is available only if the Central Admin Server feature (CAS) is installed. See “About the Central Admin Server feature” on page 1321.

5 Click **OK**.

See [“Backup networks”](#) on page 716.

Running the next scheduled backup job before its scheduled time

You can run the next scheduled backup job in a backup definition at any time. You may want to run a scheduled backup job early to ensure that important data gets backed up or to make sure that a scheduled job completes successfully. Running a scheduled backup job early does not affect its regular schedule. The job still runs normally as scheduled.

To run the next scheduled backup job

- 1 On the **Backup and Restore** tab, do one of the following:
 - To run the next scheduled backup for a single server's backup jobs, right-click the server name.

- To run the next scheduled backup for multiple servers' backup jobs, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Click **Run Next Backup Now**.
 - 3 Click **Yes** to confirm that you want to run the job or jobs now.
- See [“Backing up data”](#) on page 158.

Editing backup definitions

You can edit existing backup definitions. You can modify any existing backup definition's backup selections or the backup job settings for any backup jobs that it contains. The backup selections include any servers, volumes, or data that you have selected to back up. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example.

You can also add a stage to an existing backup definition to duplicate your backup data or create virtual machines with your backup data.

See [“Adding a stage to a backup definition”](#) on page 220.

If you choose to edit more than one backup definition at once, you can edit only the properties that the definitions have in common. For example, if you choose to edit two backup definitions at once and the definitions use different schedules, you cannot edit the schedules. If you do not see the settings that you want to edit, repeat this procedure, but select only one definition to edit at a time.

Note: You cannot edit a backup definition while one of its backup jobs is running.

This topic includes the following procedures:

[To edit a backup definition's backup selections or select user shares to back up](#)

[To edit a backup definition's job settings](#)

To edit a backup definition's backup selections or select user shares to back up

- 1 Do one of the following:

To edit backups from the **Backup and Restore** tab

Complete the following steps:

- On the **Backup and Restore** tab, do one of the following:
 - To edit backups for a single server, right-click the server name.
 - To edit backups for multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- Click **Edit Backups**.

Note: If the server or servers that you selected have multiple backup definitions, select the definitions that you want to edit on the **Backup Job Selection** dialog box and then click **OK**.

To edit backups from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, do one of the following:
 - To edit a single backup, right-click the job name.
 - To edit multiple backups at once, Shift + click or Ctrl + click the job names, and then right-click one of the selected jobs.
- Click **Edit**.

To edit backups from the **Storage** tab


Complete the following steps:

- On the **Storage** tab, double-click the storage device or storage pool for the backup that you want to edit.
- On the **Job History** pane, do one of the following:
 - To edit a single backup, right-click the backup.
 - To edit multiple backups, Shift + click or Ctrl + click the backups and then right-click one of the selected backups.

Note: You can only edit any backup jobs that have previously run from the **Storage** tab.

- Click **Edit Backup**.

2 In the **Selections** box, click **Edit**.



3 To add additional servers to the backup definition:

- Click the **Add (+)** button.
- Select the server or servers that you want to add to the backup definition. You can also click **New Server** to add a new server to Backup Exec.
- Click **OK**.

- 4 Select the data that you want to back up.

Servers are selected in their entirety by default. If you do not want to back up the entire server, double-click the server name to see all of the contents. Select the check boxes for each item that you want to back up.

Note: Deselecting a server's critical system components creates backup sets that cannot be used for some restore scenarios.

See [“About backing up critical system components”](#) on page 184.

- 5 Click **OK**.
- 6 When you are finished editing the backup definition, click **OK** on the **Backup Properties** dialog box.

To edit a backup definition's job settings

- 1 Do one of the following:

To edit backups from the **Backup and Restore** tab

Complete the following steps:

- On the **Backup and Restore** tab, do one of the following:
 - To edit backups for a single server, right-click the server name.
 - To edit backups for multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- Click **Edit Backups**.

Note: If the server or servers that you selected have multiple backup definitions, select the definitions that you want to edit on the **Backup Job Selection** dialog box and then click **OK**.

To edit backups from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, do one of the following:
 - To edit a single backup, right-click the job name.
 - To edit multiple backups at once, Shift + click or Ctrl + click the job names, and then right-click one of the selected jobs.
- Click **Edit**.

To edit backups from the **Storage** tab


Complete the following steps:

- On the **Storage** tab, double-click the storage device or storage pool for the backup that you want to edit.
- On the **Job History** pane, do one of the following:
 - To edit a single backup, right-click the backup.
 - To edit multiple backups, Shift + click or Ctrl + click the backups and then right-click one of the selected backups.

Note: You can only edit any backup jobs that have previously run from the **Storage** tab.

- Click **Edit Backup**.

2 In the **Backup** box, click **Edit**.



- 3 In the left pane, select the backup job setting that you want to modify.
The options vary depending on what you selected to back up.
- 4 Make any necessary changes.
- 5 When you are finished modifying any applicable options for the backup, click **OK**.
- 6 When you are finished editing the backup definition, click **OK** on the **Backup Properties** dialog box.

See [“Backing up data”](#) on page 158.

Testing credentials for backup sources

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

It is recommended that you test to make sure that you have the appropriate credentials to access the content that you want to back up. If a credentials test fails, you can enter new credentials for the content so that Backup Exec can access it.

You can test the credentials that are associated with a backup source at any time from the **Credentials** pane. You can also test the credentials when you create a backup definition.

Note: You cannot test credentials for virtual machines, but the job runs if you provide the correct credentials. If the job fails, you may need to retry the job with different credentials.

Complete any of the following procedures to test credentials for backup sources:

[To test the credentials that are associated with a backup source](#)

[To test the credentials for all of the backup sources in a specific backup definition](#)

[To test the credentials for a specific backup source in a specific backup definition](#)

To test the credentials that are associated with a backup source

- 1 On the **Backup and Restore** tab, double-click the server whose credentials you want to test.
- 2 In the left pane, click **Credentials**.
- 3 Select the backup source that you want to test.
- 4 In the **Credentials** group, click **Test Credentials**.

The **Credential Status** field displays the results of the test.

To test the credentials for all of the backup sources in a specific backup definition

- 1 Create a new backup or edit an existing backup.
- 2 In the **Selections** box, click **Test/Edit Credentials**.
- 3 Click **Test All**.

The **Credential Status** field displays the results of the test.

- 4 Click **OK**.

To test the credentials for a specific backup source in a specific backup definition

- 1 Create a new backup or edit an existing backup.
- 2 In the **Selections** box, click **Test/Edit Credentials**.
- 3 Select the backup source.
- 4 Click **Test Selected**.

The **Credential Status** field displays the results of the test.

- 5 Click **OK**.

See [“Replacing the credentials for a backup source”](#) on page 213.

See [“Creating new credentials for a backup source”](#) on page 214.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 215.

Replacing the credentials for a backup source

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

It is recommended that you test to make sure that you have the appropriate credentials to access the content that you want to back up. If a credentials test fails, you can enter new credentials for the content so that Backup Exec can access it. If you need to change the credentials that are associated with a backup source, you can do so at any time on the **Backup and Restore** tab from the **Credentials** pane. You can also change a backup source's credentials when you create or edit backups.

To replace the credentials for a backup source

- ◆ Do one of the following:

To replace the logon account that is associated with a backup source on the **Credentials** pane

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server whose credentials you want to view.
- In the left pane, click **Credentials**.
- In the **Logon Account** field, select the logon account that you want to use for the backup source.
- Click **Apply**.

To replace the logon account that is associated with a backup source in a backup definition

Complete the following steps:

- Create a new backup or edit an existing backup.
- In the **Selections** box, click **Test/Edit Credentials**.
- In the **Logon Account** field, select the logon account that you want to use for the backup source.
- Click **OK**.

See [“Testing credentials for backup sources”](#) on page 212.

See [“Creating new credentials for a backup source”](#) on page 214.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 215.

Creating new credentials for a backup source

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

If the credentials for a backup source change, you must enter the new credentials in Backup Exec and associate them with the backup source. You can create new credentials for a backup source at any time on the **Credentials** pane. You can also create new credentials for a backup source when you create or edit backup definitions.

To create new credentials for a backup source

- 1 Do one of the following:

To create new credentials for a backup source on the **Credentials** pane

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server for which you want to create new credentials.
- In the left pane, click **Credentials**.

To create new credentials for a backup source in a backup definition

Complete the following steps:

- Create a new backup or edit an existing backup.
- In the **Selections** box, click **Test/Edit Credentials**.

- 2 In the **Logon Account** field next to the backup source, select **<new logon account>**.

- 3 In the **User name** field, type the user name for the new logon account.
- 4 In the **Password** field, type the password for the new logon account.
- 5 In the **Confirm password** field, type the password again to confirm it.
- 6 In the **Account name** field, type a unique name for the new logon account.
- 7 In the **Notes** field, type any optional notes to explain how the Backup Exec logon account is used.
- 8 Select **This is a restricted logon account** if you want the Backup Exec logon account to be used only by the owner of the logon account and those who know the password.

If this option is not selected, the Backup Exec logon account is created as a common account. Common accounts are the shared accounts that all users can access.
- 9 Select **This is my default account** to make this account your default Backup Exec logon account, which is used to browse, make selections, or restore data on your local computers and remote computers
- 10 Click **OK**.

See [“Testing credentials for backup sources”](#) on page 212.

See [“Replacing the credentials for a backup source”](#) on page 213.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 215.

Deleting retired or unused backup sources from the Credentials pane

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If you remove one of the backup sources from your environment, it still displays in the list of backup sources on the **Credentials** pane. You can delete retired or unused backup sources from the pane to help make it easier to manage. Backup Exec browses and discovers resources during normal operations. If the backup source that you delete still exists in your environment, Backup Exec will add it to the list of resources again the next time that it is discovered. You should only delete backup sources from the **Credentials** pane if they no longer exist in your environment.

To delete retired or unused backup sources from the **Credentials** pane

- 1 On the **Backup and Restore** tab, double-click the server that contains the backup source that you want to delete from the **Credentials** pane.
- 2 In the left pane, click **Credentials**.
- 3 Select the backup source that you want to delete from the **Credentials** pane.
- 4 In the **Credentials** group, click **Delete Selected Resource**.

Backup Exec removes the backup source from the list.

Note: If you try to delete a backup source that is used in a backup definition, Backup Exec does not delete that backup source. When you delete a backup source that still exists in your environment, Backup Exec repopulates the list with the backup source the next time it browses your environment.

See [“Testing credentials for backup sources”](#) on page 212.

See [“Replacing the credentials for a backup source”](#) on page 213.

See [“Creating new credentials for a backup source”](#) on page 214.

How job scheduling works in Backup Exec

Backup Exec lets you configure the time and the frequency for which you want to run backup jobs. You can run jobs immediately, once on a specific day and time, or more than once according to a schedule. Backup Exec lets you use minutes, hours, days, weeks, months, or years as measurements of time to create a recurring pattern for the schedule. Alternatively, you can select specific days of the month to create a recurring schedule on which jobs should run.

This topic includes the following information:

[Availability window](#)

[Scheduling conflicts](#)

[Including and excluding dates](#)

Availability window

Backup Exec has additional scheduling options that you can use to ensure that the job does not run outside of your availability window. The availability window is the time range when your backup sources are available to be backed up. You can configure how long you want an incomplete job to remain scheduled before Backup Exec reschedules the job and changes its completion status to Missed. You can also configure the job to automatically be canceled if it is running for

too long after its scheduled start time. These options help ensure that backup jobs do not affect your system resources during critical hours.

See [“Configuring schedules for backup jobs”](#) on page 644.

Scheduling conflicts

Backup Exec resolves scheduling conflicts between two backup jobs by running the less common job and skipping the more common job. If a full backup job and an incremental or differential backup job are scheduled to run at the same time, Backup Exec runs the full backup. The incremental backup or differential backup is skipped and it runs again at its next scheduled time. Full backup jobs always supersede incremental and differential backup jobs. If two backup jobs of the same type are scheduled to run at the same time, Backup Exec runs the less frequently scheduled job. The more frequently scheduled job is skipped. The skipped backup job runs again at its next scheduled time. For example, if a monthly job and a daily job are scheduled to run at the same time, Backup Exec runs the monthly job. The daily job is skipped and it runs on the following day per its normal schedule.

Including and excluding dates

The **Exclude Dates** option lets you exclude specific dates from the schedule. For example, you can exclude holidays from your normal backup schedule.

You can use the **Include Dates** option to include dates with your backup schedule. When you include a date with your backup schedule, the backup job runs on the included date, even if it is not normally scheduled to run on that date. You may want to schedule an include date to run an extra backup job, outside of the job's normal schedule.

See [“Preventing backup jobs from running on a specific date”](#) on page 218.

See [“Including a specific date in the schedule for a backup job”](#) on page 217.

Including a specific date in the schedule for a backup job

You can include specific dates with your backup schedule for a backup job. When you include a date with your backup schedule, the backup job runs on the included date, even if it is not normally scheduled to run on that date. You may want to schedule an include date to run an extra backup job, outside of the job's normal schedule. Scheduling an include date does not affect a backup job's normal schedule.

Note: Included dates are applied to the job schedule before excluded dates. When any excluded dates are applied, they may overwrite the included dates if there are any conflicts. Therefore, if you select the same date as both an included date and an excluded date, Backup Exec excludes the date from your backup schedule.

See [“Preventing backup jobs from running on a specific date”](#) on page 218.

To include a specific date in the schedule for a backup job

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 Click the **Recurrence** field for the job to which you want to include a date.
- 5 On the **Include/exclude dates** tab, click **Include dates**.
- 6 Do one of the following:

To manually enter the date Complete the following steps:

- In the **Select Date** field, type the date that you want to include with the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar Click the date that you want to include with the backup schedule.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

- 7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 216.

Preventing backup jobs from running on a specific date

You can exclude specific dates, such as holidays, from your backup schedule for individual backup jobs.

When you exclude dates for a specific backup job, only that backup job is affected by the exclude date. The job does not run on the excluded date, even if it is normally scheduled to run. After the exclude date, the job resumes running on its normal schedule.

Note: Included dates are applied to the job schedule before excluded dates. When any excluded dates are applied, they may overwrite the included dates if there are any conflicts. Therefore, if you select the same date as both an included date and an excluded date, Backup Exec excludes the date from your backup schedule.

See [“Including a specific date in the schedule for a backup job”](#) on page 217.

To prevent backup jobs from running on a specific date

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 Click the **Recurrence** field for the job to which you want to add an exclude date.
- 5 On the **Include/exclude dates** tab, click **Exclude dates**.
- 6 Do one of the following:

To manually enter the date Complete the following steps:

- In the **Select Date** field, type the date that you want to exclude from the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar Click the date that you want to exclude.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

- 7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 216.

Viewing all scheduled backup jobs on a calendar

You can view all of your scheduled backup jobs for a month, for a week, or for a day on a calendar. It can be helpful to view your backup jobs in the calendar format to make sure there are no scheduling conflicts. You may want to check the calendar before you create a new job.

To view all scheduled backup jobs on a calendar

- 1 On the **Backup and Restore** tab, in the **Backups** group, click **Backup Calendar**.
- 2 When you are finished viewing the calendar, click **Close**.

See [“How job scheduling works in Backup Exec”](#) on page 216.

Adding a stage to a backup definition

Stages are the additional tasks that you can run with backup jobs as part of the backup definition. You may choose to add stages to the backup definition to customize it. You can add one or more stages for virtualization and duplication when you create a backup definition. Stages can also be added to existing backup definitions.

For example, you may create a backup job that backs up any important data that must be sent off-site. You can add a duplicate stage to the backup definition that contains that job. The duplicate stage automatically sends the backup data to tape storage when the backup job is complete. Then you can take the tape off-site to ensure that your data is safe.

Table 4-8 Types of stages

Stage	Description
Duplicate to Disk	Creates a duplicate copy of your backup and sends it to disk storage.
Duplicate to Tape	Creates a duplicate copy of your backup and sends it to tape storage.
Duplicate to Cloud	Creates a duplicate copy of your backup and sends it to the cloud storage.
Convert to Virtual Machine After Backup	Creates a virtual machine from your backup sets after the backup job is complete. See “How conversion of physical computers to virtual machines works in Backup Exec” on page 589.

Table 4-8 Types of stages (*continued*)

Stage	Description
Convert to Virtual Machine Simultaneously with Backups	Creates a virtual machine from your backup sets while the backup job is running. See “How conversion of physical computers to virtual machines works in Backup Exec” on page 589.

To add a stage to a backup definition

- 1 Create or edit a backup definition.

Note: You cannot add a stage to one-time backups.

- 2 In the **Backup** box, click **Add Stage**.
- 3 Select the type of stage that you want to add.
You can add multiple stages to a backup definition.
- 4 In the stage box, click **Edit**.
- 5 In the left pane, click **Schedule**, and then select when you want the stage to run.
- 6 In the left pane, click **Storage**, and then select the storage device that you want to use for this stage.
- 7 Select any additional options that apply to this stage.
- 8 When you are finished selecting all applicable options for this stage, click **OK**.

Repeat this procedure for each stage that you want to add to the backup definition.

See [“Backing up data”](#) on page 158.

See [“Editing a stage”](#) on page 221.

Editing a stage

You can edit a stage that is part of a backup definition.

To edit a stage

- 1 Do one of the following:

To edit a stage from the **Backup and Restore** tab

Complete the following steps:

- On the **Backup and Restore** tab, right-click the server that contains the backup definition with the stage that you want to edit.
- Click **Edit Backups**.
- If the server has more than one backup definition, select the definition that contains the stage that you want to edit, and then click **OK**.

To edit a stage from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, right-click the job that contains the stage that you want to edit.
- Click **Edit**.

- 2 Click **Edit** in the box that contains the stage that you want to edit.
- 3 Make any necessary changes.
- 4 When you are finished making changes to the stage, click **OK** on the **Backup Properties** dialog box.

See [“Adding a stage to a backup definition”](#) on page 220.

Duplicating backup sets or a job history manually

You can configure a duplicate stage to automatically duplicate backup data after a backup job runs. You may want to duplicate data to have an extra copy to send off-site, for example.

See [“Adding a stage to a backup definition”](#) on page 220.

You can also manually duplicate backup data from completed jobs at any time. When you duplicate backed up data from completed jobs, you select the backup sets or job history that you want to duplicate. The data that you select is read from the source and written to the selected destination, such as a drive, drive pool, or backup folder. You can encrypt the duplicated data. You can schedule when this type of job runs, but it only runs one time.

You can select to duplicate one or more individual backup sets or you can duplicate an entire job history. You should duplicate backup sets if you want to duplicate only the data that was backed up in a specific backup job instance. When you duplicate a job history, Backup Exec includes all of a backup definition's dependent backup sets. For example, if you select to duplicate an incremental backup job, Backup Exec automatically duplicates all incrementals dating back to, and including, the last full backup job.

This topic includes the following information:

[Duplicating data from a virtual device to a physical device](#)

[To duplicate backup sets](#)

[To duplicate job history](#)

Duplicating data from a virtual device to a physical device

You can use a duplicate backup job to copy data directly from a virtual device to a physical device. Software encryption cannot be applied to a duplicate backup job when you copy data directly from a virtual device to a physical device. You must either disable DirectCopy or select not to encrypt the job.

See [“Copying data from a virtual tape library to a physical tape device using DirectCopy to tape”](#) on page 229.

If you duplicate any Oracle backup sets that were created with multiple data streams, note the following:

- Backup Exec converts the multiple data streams to a sequential data stream during the duplication job.
- A restore job from the duplicated copy may be slower than a restore job from the original media.

To duplicate backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to duplicate.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To duplicate a single backup set, right-click the backup set.
 - To duplicate multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Select **Duplicate**.

- 5 In the **Schedule** group box, select when you want Backup Exec to run the duplicate job:

To run the duplicate job immediately	Select Run now .
To schedule the job for a later time	Select Run on , and then enter the date and time.
To schedule the job to run later using an external scheduling tool	Select Create without a schedule . You can schedule the job to run later using an external scheduling tool.
To submit the job on hold	Click Submit job on hold . The job is created with an On Hold status. It remains on hold until you remove the hold on it.

- 6 In the **Storage** field, select the storage device to which you want to copy the backup sets.

- 7 In the **Keep for** field, select the amount of time that you want Backup Exec to keep the backup sets.

The media is protected from being overwritten for the amount of time that you specify.

Select **Use source retention** to keep the backup sets for the same amount of time as the source backup sets that you want to duplicate.

- 8 Do any of the following:

To enable compression for the duplicate backup sets	In the Compression field, select the type of compression.
To enable encryption for the duplicate backup sets	Complete the following steps: ■ In the Encryption type field, select the type of encryption. ■ In the Encryption key field, select the encryption key that you want to use or select Manage keys to create a new key.
To run a verify operation on the duplicate backup sets	Select Verify at the end of the job .

- 9 On the **Duplicate Job** dialog box, click **OK**.

To duplicate job history

1 Do one of the following:

To duplicate job history from the Backup and Restore tab or the Storage tab	Complete the following steps: ■ On the Backup and Restore tab or the Storage tab, double-click the server or the storage device that is related to the job history that you want to duplicate. ■ In the left pane, click Job History. ■ Do one of the following: ■ To duplicate a single job history, right-click the job history. ■ To duplicate multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

To duplicate job history from the Job Monitor tab	Complete the following steps: ■ On the Job Monitor tab, do one of the following: ■ To duplicate a single job history, right-click the job history. ■ To duplicate multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

2 Click **Duplicate**.

3 In the **Schedule** group box, select when you want Backup Exec to run the duplicate job:

To run the duplicate job immediately	Select Run now .
To schedule the job for a later time	Select Run on , and then enter the date and time.
To schedule the job to run later using an external scheduling tool	Select Create without a schedule . You can schedule the job to run later using an external scheduling tool.
To submit the job on hold	Click Submit job on hold . The job is created with an On Hold status. It remains on hold until you remove the hold on it.

- 4 In the **Storage** field, select the storage device to which you want to copy the job history.
- 5 In the **Keep for** field, select the amount of time that you want Backup Exec to keep the job history.

The media is protected from being overwritten for the amount of time that you specify.

Select **Use source retention** to keep the job history for the same amount of time as the source job history that you want to duplicate.

- 6 Do any of the following:

To enable compression for the duplicate job history In the **Compression** field, select the type of compression.

To enable encryption for the duplicate job history Complete the following steps:

- In the **Encryption type** field, select the type of encryption.
- In the **Encryption key** field, select the encryption key that you want to use or select **Manage keys** to create a new key.

To run a verify operation on the duplicate job history Select **Verify at the end of the job**.

- 7 On the **Duplicate Job** dialog box, click **OK** for each duplicate job that you create.

Running a test run job manually

Test run jobs attempt to determine if a scheduled backup could possibly fail when you run it. When you run a test job, no data is backed up. Instead, Backup Exec checks your storage capacity, credentials, and media to find potential errors. If there is an error, the job continues to run until it is completed. The error appears in the job log. You can also configure Backup Exec to send a notification to a designated recipient.

During a test run job, the following things may cause a job to fail:

- Logon credentials are incorrect.
- Storage capacity is not sufficient.
- Tape cartridge media or disk cartridge media is not available.
- Overwritable media is not available for an overwrite job.

- Appendable media is not available for an append job.

A test run job checks the media capacity that is available for the selected job. However, you can check if there is enough available media for multiple test run jobs in the Test Run Results Report.

See [“Test Run Results report”](#) on page 839.

You can configure test run jobs to run automatically before your scheduled backup jobs. Or you can manually run a test run job at any time.

Before you run a test run job, it is recommended that you run backup jobs to your storage devices first. Backup Exec does not recognize the capacity of a storage device until an actual backup job sends data to the device. If you create a test run job before any other jobs, Backup Exec cannot check that the device has sufficient capacity to perform the backup job. After at least one backup job has sent data to a device, Backup Exec can determine the capacity.

To run a test job manually

- 1 Do one of the following:

To run a test job from the
Backup and Restore tab

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server that contains the job you want to test.
- In the left pane, click **Jobs**.

To run a test job from the
Job Monitor tab

Select the **Job Monitor** tab.

- 2 Right-click the job that you want to test, and then click **Test Run**.
- 3 Click **Yes** to confirm that you want to run the test job now.

See [“Configuring automatic test run jobs for backup jobs”](#) on page 659.

Verifying backed up data manually

Backup Exec can perform a verify operation to make sure that the media can be read after a backup job has been completed. It is recommended that you verify all backed up data to ensure the integrity of the collection of data and the media on which it resides. By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the verify operation to take place at a later time or disable the verify operation altogether. You can change Backup Exec's verify options as part of the default backup settings or for individual backup jobs.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

You can also choose to manually run a verify operation on a backup set or a job history at any time. You can verify backup sets if you want to verify only the data that was backed up in a specific backup job instance. If you want to verify a backup definition and all of its dependent backup sets, you can verify a job history. For example, if you want to verify a backup definition that used incremental backups, Backup Exec verifies all incrementals dating back to, and including, the last full backup.

This topic includes the following procedures:

[To verify specific backup sets](#)

[To verify a job history](#)

To verify specific backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup set or backup sets that you want to verify.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To verify a single backup set, right-click the backup set.
 - To verify multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Verify**.
- 5 In the **Schedule** group box, select when you want Backup Exec to run the verify operation:

To run the verify operation immediately Select **Run now**.

To schedule the operation for a later time Select **Run on**, and then enter the date and time.

To schedule the job to run later using an external scheduling tool Select **Create without a schedule**.
 You can schedule the verify operation to run later using an external scheduling tool.

- 6 Click **OK**.

To verify a job history

- 1 Do one of the following:

To verify a job history from the **Backup and Restore** tab or the **Storage** tab

Complete the following steps:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the job history that you want to verify.
- In the left pane, click **Job History**.
- Do one of the following:
 - To verify a single job history, right-click the job history.
 - To verify multiple job histories, Shift + click or Ctrl + click the job history, and then right-click one of the selected job histories.

To verify a job history from the **Job Monitor** tab

On the **Job Monitor** tab, do one of the following:

- To verify a single job history, right-click the job history.
- To verify multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

2 Click **Verify**.

3 In the **Schedule** group box, select when you want Backup Exec to run the verify operation:

To run the verify operation immediately

Select **Run now**.

To schedule the operation for a later time

Select **Run on**, and then enter the date and time.

To schedule the job to run later using an external scheduling tool

Select **Create without a schedule**.

You can schedule the verify operation to run later using an external scheduling tool.

4 Click **OK**.

Copying data from a virtual tape library to a physical tape device using DirectCopy to tape

Backup Exec's **DirectCopy to tape** option enables data to be copied from a virtual tape library directly to a physical tape device during a duplicate backup job. The Backup Exec server coordinates the duplicate job, but it does not copy the data.

Instead, the virtual tape library copies the virtual tape image directly to the physical device. The Backup Exec server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual tape library or the physical device. The job log for the duplicate backup job indicates that **DirectCopy to tape** is enabled.

To use **DirectCopy to tape**, both the source device and the destination device must be NDMP-enabled. If the devices are not NDMP-enabled, then Backup Exec performs a regular duplicate backup job.

Both hardware encryption and software encryption are supported with **DirectCopy to tape**. For software encryption, both the source backup set and the destination backup set must use software encryption.

Note: If you select disk storage as the destination device for a duplicate job with **DirectCopy to tape** enabled, Backup Exec performs a regular duplicate job.

Table 4-9 How to use DirectCopy to tape to copy data from a virtual tape library to a physical device

Step	Notes	For more information
Create a regular backup job.	Select a virtual tape library as the storage destination.	See “Backing up data” on page 158. See “Configuring storage options for backup jobs” on page 652.
Create a duplicate backup job.	In the DBA-initiated job settings: ■ Select a physical tape device as the destination. ■ Select Enable DirectCopy to tape.	See “DBA-initiated job templates” on page 746. See “Storage options for DBA-initiated jobs” on page 749. See “Duplicate job settings for DBA-initiated jobs” on page 755.

About using a self-protect backup job to protect server metadata

Backup Exec metadata contains information about backups and media server configuration. This metadata includes records of the files that are backed up, the media on which the files are stored, all media, and storage devices information. If the metadata is not protected, you may lose information on the media server.

Backup Exec has now introduced a self-protect backup job that is automatically created when Backup Exec is installed and performs backup of the metadata on the media server. The format of the job is as follows: <media-server-name>

Backup Exec Self Protect Job.

The backup selection list cannot be edited as the list is pre-defined. By default, this job is targeted to the **Any Storage** option and it is recommended to configure a dedicated storage for the backup job. You can update the backup job target storage and schedule based on your preference.

The self-protect backup sets contain components such as catalog files, database, database encryption keys, job logs and so on, which can be restored to any existing media server. It is recommended to enable encryption to protect backup data from any unauthorized access.

The job cannot be deleted because it protects the essential metadata of a server. It can be put on hold if you are performing a media server backup through other jobs. If any metadata paths are updated, a self-protect incremental backup is converted to a full backup. For example, the catalog or job log path is updated.

In a CAS-MBES scenario, the job is automatically created on each media server and backs up metadata specific to that media server. If you want to restore from self-protect backup sets, you must perform the restore on each media server separately.

Prerequisites for self-service restore

Before you start restore from self-protect backup sets, ensure that the following requirements and prerequisites are met:

- Ensure that you have details of the current media server hostname, data path, catalog path, and job log path.
- Ensure that the new media server name (NetBIOS) and domain name are same as the source server because self-protect backup sets restore database and catalogs to the original server name path.
- Ensure that the catalog and job log path are same as the source server during Backup Exec installation on the new server because self-protect backup sets restore data, catalog and job logs files to its original server path.
- Ensure that you import the devices to the destination server (new media server) if the source media server has a deduplication storage or a cloud storage device configured.
- Ensure that you attach the same disks to the new media server. If the source media server has disks configured on the physical volume and cannot be moved, copy the BEData and the BEControl folder at the root of the volume on the target server.

Restoring Backup Exec from the self-protect backup set

To restore metadata from the self-protect backup set

- 1 Install Backup Exec with the same server name, data, catalog path, and job log path as the source server.
- 2 Import the existing storage containing self-protect backup sets to the new Backup Exec server.

See [“Importing media to Backup Exec ”](#) on page 568.

- 3 Right-click the disk storage device, and then click **Inventory and Catalog**.

See [“Inventorying and cataloging a storage device”](#) on page 563.

The self-protect backup sets are now available for restore.

- 4 On the **Backup and Restore** tab, right-click the media server, and then click **Restore**.
- 5 Select Files, folders, or volumes, and then click **Next**.
- 6 Select **File and folder backups to a point-in-time**, and then click **Next**.
- 7 Select the self-protect backup set.
- 8 Follow the Restore Wizard prompts to restore the data.

After the restore job is completed, you must restore the Backup Exec database using BEUtility.

To restore the database using BEUtility

- 1 From the Backup Exec installation directory, in <Backup Exec install path>\Backup Exec, double-click **BEUtility**.
- 2 On the Backup Exec Utility, click **All Backup Exec Servers** > right-click the media server name > click **Restore Database from BE Self Protect Backup**.
- 3 Review the information and click **OK**.

- 4 Enter the credentials that you provided when you installed Backup Exec.

These credentials are later used as Backup Exec Service Account, System Logon Account, and Owner of System Logon Account credentials.

- 5 Complete the restore process and then click **Close**.

The new Backup Exec server is now ready to use and in the same state as the original server.

The role-based access feature is disabled. Go to Backup Exec settings to enable the feature.

If any errors are observed during service start, open the service manager and restart all Backup Exec and Deduplication services.

In a CAS-MBES scenario, you must perform all the earlier steps for each Backup Exec server separately using the self-protect backup set.

Restores

This chapter includes the following topics:

- [Methods for restoring data in Backup Exec](#)
- [Searching for data to restore](#)
- [Configuring malware scan for virtual machine backup sets before restore](#)
- [Restoring data from a server, a backup set, a backup job, or a storage device](#)
- [Restoring file system data](#)
- [Performing a complete online restore of a Microsoft Windows computer](#)
- [Restoring System State](#)
- [Installing a new Windows Server domain controller into an existing domain by using a redirected restore](#)
- [Restoring Backup Exec Shadow Copy Components](#)
- [Restoring utility partitions or Unified Extensible Firmware Interface system partitions](#)
- [About restoring encrypted data](#)
- [About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec](#)
- [Canceling a restore job](#)
- [How Backup Exec catalogs work](#)
- [Cataloging backup sets](#)

Methods for restoring data in Backup Exec

Backup Exec provides guided **Search** and **Restore** methods to assist you when you search for or restore backed up data.

From **Search** or **Restore** on the **Backup and Restore** tab, you can do the following:

- Restore data to the location from which it was originally backed up or redirect the restore to another location.
- Start the restore job immediately or schedule it to run at a future time.

Note: The **Restore Wizard** displays only up to 30,000 items. If you want to restore data from a folder that contains more than 30,000 items, you should search for the item that you want to restore. You can use search criteria such as the type of data and the date of the original backup to help reduce the number of items that displays.

Table 5-1 Guided methods to search for and restore data

Method	Description
Search	Lets you select multiple servers on which to search for backup sets. Then, you can choose to restore the data, or you can copy and save the search criteria and the results to the clipboard. You can then email the results to the person who requested the restore to ensure that you have found the correct data before you restore it. To start the Search Wizard, on the Backup and Restore tab, select one or more servers, and then in the Restores group, click Search. See “Searching for data to restore” on page 237.

Table 5-1 Guided methods to search for and restore data (continued)

Method	Description
Restore	Lets you browse the backup sets from a single server, and then restore the data. You can restore file system data, System State data, Backup Exec Shadow Copy Components, utility partitions or UEFI system partitions, and more. You can also perform a complete online restore of a Windows computer if the computer was fully selected for a backup. By default, backup jobs include all necessary components that are required for a complete restore. To start the Restore Wizard, on the Backup and Restore tab, select a server, and then in the Restores group, click Restore. See “Restoring data from a server, a backup set, a backup job, or a storage device” on page 241. See “Performing a complete online restore of a Microsoft Windows computer” on page 246.
Simplified Disaster Recovery	Lets you recover Windows computers after a hard drive failure. The Simplified Disaster Recovery wizards guide you in preparing for disaster recovery, and in recovering a local computer or a remote computer to its pre-disaster state. See “About Simplified Disaster Recovery” on page 877.

- See [“Restoring file system data”](#) on page 244.
- See [“Restoring System State ”](#) on page 247.
- See [“Restoring Backup Exec Shadow Copy Components”](#) on page 251.
- See [“Restoring utility partitions or Unified Extensible Firmware Interface system partitions”](#) on page 252.
- See [“Restoring Exchange data”](#) on page 1196.
- See [“Restoring SQL databases and transaction logs”](#) on page 1146.
- See [“About restoring Oracle resources”](#) on page 1248.
- See [“Restoring VMware virtual machines and vmdk files”](#) on page 1034.
- See [“Restoring a deduplication disk storage device or deduplicated data”](#) on page 1001.
- See [“Restoring Microsoft SharePoint data”](#) on page 1214.
- See [“Restoring Enterprise Vault”](#) on page 1277.

See [“About restoring data to Linux and Unix computers”](#) on page 1441.

Searching for data to restore

You can select one server or multiple servers on which to search for backup sets. Then, you can choose to restore the data, or you can copy and save the search criteria and the results to the clipboard. You can then email the results to the person who requested the restore to ensure that you have found the correct data before you restore it.

Backup Exec creates separate restore jobs for each server that you restore data to.

The **Search Wizard** supports only the following types of data:

- Files and folders
- Exchange and SharePoint backup sets for which Granular Recovery Technology was enabled

Note: You can search the backup sets for any data that was backed up from a virtual machine if you have selected one of the full catalog options on the **Instant GRT and Full Catalog Options** page and the full catalog job is complete. However, the search option is available in the **Restore Wizard** only when you select an application installed on the virtual machine for restore.

Search is not available if you have selected the **Enable Instant GRT** option on the **Instant GRT and Full Catalog Options** page.

For files and folders, if you have blocked access to backed up files using the `Import-BEItemsToBlock` BEMCLI command, the blocked files are not available when you are browsing for files to be restored.

To search for data to restore

- 1 On the **Backup and Restore** tab, right-click the server or servers on which you want to search for data, and then click **Search**.
- 2 Follow the **Search Wizard** prompts to search for and restore data.

Configuring malware scan for virtual machine backup sets before restore

Malware Detection feature provides the ability to perform a scan on virtual machine backup sets. Malware scan operation can be run as an ad hoc job on a given backup set or it can be run as part of the restore operation.

You can use the feature for the following scenarios:

- To ensure a malware free backup image before a restore operation for business continuity.
- To perform a scan before sending a copy of the backup set to a secondary storage.
- To scan specific backups of interest for compliance purposes, recovery readiness, or any other reason.

The scan can be run on VMware and Hyper-V virtual machines and only supports GRT enabled backup sets for disk and deduplication disk devices.

All the previous backup sets in the backup chain of the selected set should be GRT enabled backup sets on supported devices.

Backup Exec now supports Symantec Protection Engine (SPE) that can be used for malware scan. The malware scan settings dialog now provides the option to select the scan engine that can be used to perform malware scan. You can choose either Microsoft Windows Defender or SPE. By default, Backup Exec uses Windows Defender to run the malware scan on the backup set chain. Ensure that you have enabled Microsoft Defender and have the updated antivirus definitions. Backup Exec depends on the output of Windows Defender to determine if a backup set is malware free. Ensure that you have valid license for the SPE and the SPE services are running when you choose SPE for malware scan. Backup Exec depends on the output of the SPE to determine if a backup set is malware free.

For an incremental backup, the previous backup sets in the recovery chain are also used for the scan, but only the selected incremental set is marked as scanned.

To perform a malware scan, the selected backup set and previous sets in its backup chain are mounted. Each volume of the virtual machine from the mounted backup set is scanned for malware using the scan engine that is selected on the Malware scan settings.

You can choose to run malware scan as part of the physical virtual machine restore or GRT restore operation. If you perform a GRT restore, the entire virtual machine backup is scanned before performing the restore.

Backup Exec lets you modify the settings of the malware scan job. Scanning a backup set is CPU intensive and can take significant time if the virtual machine backup has large data to scan. This increases the time taken by the restore job and should be scheduled accordingly. It may affect other applications if you run multiple scan jobs. Go to **Backup Exec Settings > Malware Scan** to update the settings.

Malware scan is only supported for Backup Exec media server with Windows 2016 or later. Malware scan jobs can be run on an ad hoc basis by selecting the VMware and Hyper-V virtual machine backup sets. An ad hoc malware scan job cannot be run for GRT sets of a virtual machine.

Malware scan status for virtual machines

To see the malware scan status, on the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to view and then click **Backup Sets**.

For the selected server or storage view, the **Malware Scan Status** and **Malware Scan Time** fields are displayed.

The scan time is the time when malware scan was last run on that set.

The scan operation displays the following malware status for the virtual machine backup set:

- **Malware found**
Indicates that malware is found in one or more volumes that are mounted from the backup set.
- **Malware not found**
Indicates no malware being found in the scan of all the volumes of the backup set.

Best practices for using malware scan

Note the following best practices for using malware scan:

- It is recommended to run only one malware scan job at a time.
- It is recommended to schedule a malware scan job outside the backup window.
- It is recommended to not have any other antivirus on Backup Exec if Windows Defender is enabled for malware scan.
- It is recommended to not have any other antivirus on Backup Exec if Symantec Protection Engine (SPE) is used for malware scan.
- It is recommended to turn off the Real-Time of Microsoft Windows Defender, if you are using SPE for malware scan.

- It is recommended to have the valid license and SPE services running when you are using the SPE for malware scan.
- It is recommended to keep Windows Defender virus definitions updated on the Backup Exec server when malware scan is used.
- It is recommended to use default values in job and configuration sets while running the scan job.

Limitation of malware scan

Note the following limitation:

- Malware scan cannot be run if the backup set is created using the Faster Processing method for Hyper-V.

How to run a malware scan on a backup set

To run a malware scan on a backup set

- 1 On the **Backup and Restore** tab or the Storage tab, double-click the VMware vCenter or Hyper-V server or the storage device that is related to the backup sets.
- 2 In the left pane, click **Backup Sets**.
- 3 To run malware scan for a single backup set, right-click the backup set.
- 4 In **Backup Sets**, enter the malware scan job name.
Details of the backup set and the backup set chain are displayed.
- 5 In the left pane, click **Schedule**, and then select one of the following scheduling options:

Run now	Lets you run the malware scan job immediately.
Run on	Lets you set the date and time to run the malware scan job.
Create without a schedule	Creates the malware scan job without scheduling it. When you use this option, the scan does not run at the time of creation. The scan job remains unscheduled until you choose to run it.
Submit job on hold	Submits the malware scan job, but does not run the job until you change the job's hold status.

- 6 In the left pane, select **Notification**, and then select recipients to notify when the malware scan job completes.
- 7 Click **OK**.

Restoring data from a server, a backup set, a backup job, or a storage device

You can restore data by launching the **Restore Wizard** directly from a server, from backup sets, from completed backup jobs, and from storage device media.

You can also perform a complete online restore of a Windows computer if the computer was fully selected for a backup. By default, backup jobs include all necessary components that are required for a complete restore.

For the file and folders, if you have blocked access to backed up files using the `Import-BEItemsToBlock BEMCLI` command, the blocked files are not displayed in the search results.

See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 246.

Table 5-2 Restore methods

Restore method	More information
To restore from a server	Restoring data from a server
To restore from a backup set	To restore data from a backup set
To restore from a completed backup job	Restoring data from a completed backup job
To restore from storage device media	Restoring data from storage device media

Note: If the File Server Resource Manager (FSRM) is running on the server that you are restoring, the job may fail with an "out of disk space" error. If this occurs, disable RSRM, and then run the job again.

Restoring data from a server

You can browse the backup sets from a single server, and then restore the data.

Note: If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. The status next to the new server name indicates that it is backed up. However, you should select the icon with the old server name to restore any data that you backed up before you changed the server name.

To restore data from a server

- 1 On the **Backup and Restore** tab, in the list of servers, or on the **Job Monitor** tab, right-click a server that has been backed up.
- 2 Click **Restore**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from a completed backup job

You can restore data from a completed backup job. When you restore data from backup jobs, you can only choose data that is contained in the backup sets that the backup job produced.

To restore data from a completed backup job

- 1 Do one of the following:

To restore data from the **Backup and Restore** tab

- On the **Backup and Restore** tab, in the list of servers, double-click a server that has been backed up.
- In the **Jobs** view, expand a backup definition.
- Under the backup definition, right-click the backup job from which you want to restore data.
- Select **Restore backup sets created by this job**.

To restore data from the **Job Monitor** tab

- On the **Job Monitor** tab, in the **Jobs** pane, right-click a server that has been backed up.
- Select **Restore backup sets created by this job**.

- 2 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from a backup set

You can restore data from the backup sets that appear on the **Backup and Restore** tab. When you restore data from backup sets, you can only choose data that is contained in the backup sets.

To restore data from a backup set

- 1 On the **Backup and Restore** tab, in the list of servers, double-click a server that has been backed up.
- 2 In the **Backup Sets** view, expand a backup set from which you want to restore.
- 3 Right-click the backup set from which you want to restore data

- 4 Click **Restore**.
- 5 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from storage device media

You can restore data from the media that are contained in the storage devices that appear on the **Storage** tab. When you restore data directly from storage device media, you can only restore data that is contained in the backup sets on the media. The Restore Wizard prompts you to select options for each type of data, and then submits a separate job for each type of data.

When you restore from storage device media, you can restore multiple types of data. Separate restore jobs are submitted for each type of data.

To restore data from storage device media

- 1 Select from the following:

To restore from a disk storage device	Do the following in the order listed: ■ On the Storage tab, in the All Storage view, navigate to the storage device from which you want to restore. ■ Double-click the disk storage device, and then click Backup Sets on the left. ■ Right-click the backup set from which you want to restore, and then click Restore.
To restore from tape drive media or a disk cartridge media	Do the following in the order listed: ■ On the Storage tab, in the All Storage view, expand Tape and Disk Cartridge Media. ■ Navigate to the media that contains the data that you want to restore. ■ Right-click the media from which you want to restore, and then click Restore.

To restore from media within a robotic library Do the following in the order listed:

- On the **Storage** tab, in the **All Storage** view, navigate to the robotic library from which you want to restore.
- Expand the robotic library.
- Double-click **Slots**.
- In the Slots view, right-click the slot that contains the media from which you want to restore, and then click **Restore**.

- 2 Follow the **Restore Wizard** prompts to restore the data.
- 3 (Optional) To restore multiple types of data, proceed through the **Restore Wizard** and select the appropriate options for the type of data that you select.

Then, do the following:

- On the summary page, click **Continue** to submit the job and return to the page that lists the different types of data.
The **Job Submitted** column displays an icon to indicate that the restore job for the first type of data has been submitted.
- Select the next type of data and the appropriate restore options for that type of data. If you do not want to restore the other types of data, you can click **Cancel** to exit the **Restore Wizard** without affecting any of the jobs that you already submitted.
- After you configure jobs for each type of data, the **Continue** button changes to a **Finish** button. When you click **Finish**, Backup Exec submits the final restore job, and then closes the **Restore Wizard**.

Restoring file system data

When you restore files, folders, or volumes, you can restore to a point-in-time or you can restore from a backup set.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

To restore file system data

- 1
- On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 2
- Select **Files, folders, or volumes**, and then click **Next**.
- 3
- Do one of the following:

To restore the data to a point-in-time at which the backup set was created Note: When you restore file system data to a point-in-time, you select a backup set to restore. The backup set represents the file system data at the specific point-in-time at which it was backed up. Backup Exec automatically restores any related backup sets that are required to restore the file system data to its state at the point-in-time.	Select File and folder backups to a point-in-time.
To restore file and folder backups from a backup set Note: When you restore file system data from a backup set, you select a backup set to restore. Backup Exec restores the file system data exactly as it exists in the backup set that you select. No dependent backup sets are selected for the restore.	Select File and folder backups from a backup set.
To search for files and folders	Select Files and folders located through Search.
To restore blocked files Blocked items contain personally identifiable information. By default, the blocked items are not available for restore. Only the owner of a System Logon Account has the privileges to restore these blocked items and reason for restore is recorded in the audit log.	Do the following in the order listed: 1 Select the Allow restore of blocked items check box. 2 On the Restore Blocked Items dialog box, enter the reason for restoring blocked files and then click OK. See “About GDPR Guard” on page 841.

4 Click **Next**.

If you have blocked access to backed up files using the `Import-BEItemsToBlock BEMCLI` command, the blocked files are not available when you are searching for files to be restored.

5 Follow the **Restore Wizard** prompts to restore the data.

Performing a complete online restore of a Microsoft Windows computer

You can perform a complete online restore of a Microsoft Windows computer if the computer was fully selected for a backup. You select the backup set time from which you want to recover the computer. All required backup sets are automatically selected. You can select additional backup sets to restore as appropriate. You cannot redirect the restore of the computer.

Note: Online restore using the Agent for Windows does not restore the WindowsApps folders on operating systems that run Windows 8 or later. However, the restore job is successful. Microsoft recommends to restore the WindowsApps folders using Device Reset on the Settings panel.

The WindowsApps folders that are ignored during restore could be the following:

The folder pointed by

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Appx\PackageRoot
```

The folder pointed by

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Appx\PackageRepositoryRoot
```

```
%SystemRoot%\InfusedApps
```

To run a complete online restore of a Microsoft Windows computer

- 1 On the **Backup and Restore** tab, right-click the computer that you want to restore, and then click **Restore**.
- 2 Select **A Microsoft Windows computer that was fully selected for a backup**, and then click **Next**.
- 3 Select **Complete online restore of a computer, or restore system components**, and then click **Next**.
- 4 Follow the **Restore Wizard** prompts to restore the data.

If you have specified any blocked files, they are not restored.

See [“About backing up critical system components”](#) on page 184.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

Restoring System State

Depending on the version of Microsoft Windows, service pack levels, and features that are installed, you can restore the following system state data:

- Active Directory
- Automated system recovery
- Background Intelligent Transfer Service
- COM+ Class Registration database
- Dynamic Host Configuration Protocol
- Event logs
- File Server Resource Manager
- Internet Information Service (IIS)
- Microsoft Search Service
- Network Policy Server
- Registry
- Remote Storage
- Removable Storage Manager
- Shadow Copy Optimization Writer
- System files
- Terminal Server Licensing
- Terminal Services Gateway
- Windows Deployment Services
- Windows Management Instrumentation

Note: To restore the Windows Internet Name Service (WINS), you must also restore the registry. You cannot restore WINS by itself.

If the server is a certificate server, then System State includes the Certificate Services database.

If the server is a domain controller, then System State includes the Active Directory services database and the SYSVOL directory.

See [“About the Agent for Microsoft Active Directory”](#) on page 1311.

You must restart the computer after you restore System State data.

Warning: You should not cancel a System State restore job. Canceling this job can leave the server unusable.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

Note: When you restore the System State, Backup Exec may create the following temporary directory:

%SystemRoot%\F52E2DD5-CE7D-4e54-8766-EE08A709C28E

After the restore job finishes, you can delete the directory.

To restore System State

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore System State, and then click **Restore**.
- 2 Select **Complete online restore of a computer, or restore system components**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.
- 4 After you restore System State data, restart the computer.

See [“About backing up critical system components”](#) on page 184.

See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 246.

Restoring System State to a domain controller

To restore System State to a computer that is a domain controller, you must start the computer in safe mode. Then, use the Directory Services Restore Mode to perform the restore.

To replicate Active Directory to the other domain controllers that exist in the domain, you must perform an authoritative restore of the Active Directory. An authoritative restore ensures that the restored data is replicated to all of the

servers. Performing an authoritative restore includes running Microsoft's Ntdsutil utility after Backup Exec restores System State, but before you restart the server. For more information about authoritative restore and the Ntdsutil utility, see your Microsoft documentation.

To restore System State to a domain controller

- 1 Retart the destination server in Directory Services Restore Mode.
 See your Microsoft documentation for details on how to do this.
- 2 Open the services.
- 3 For each Backup Exec service listed, do the following in the order listed:
 - Click **Properties**.
 - On the **Log On** tab, click **This account**.
 - Enter a user account with local administrator's rights, and then click **OK**.
 - Right-click the service, and then click **Start**.
- 4 After the Backup Exec services have started, run the **Restore Wizard** to restore System State.
 See "[Methods for restoring data in Backup Exec](#)" on page 235.
- 5 In the **Restore Wizard**, enable the option **Mark this server as the primary arbitrator for replication when restoring SYSVOL in System State**.
- 6 Restart the server before you restore more data.

Installing a new Windows Server domain controller into an existing domain by using a redirected restore

To install a new Windows Server domain controller into an existing domain, the Active Directory and SYSVOL data must be replicated to the new domain controller. If there is a large amount of data to be replicated or if the connection between the domain controllers is slow, the replication time can be lengthy. The amount of data to be replicated and the connection speed also affects the Active Directory Application Mode replication time. To decrease the replication time for Active Directory and Active Directory Application Mode, you can use the Microsoft Windows feature called **Install from Media**.

For Active Directory, use the **Install from Media** feature. Restore the system state backup sets of an existing domain controller in the domain in which you want to add a new domain controller. Then, perform a redirected restore of the system state backup sets to the destination domain controller.

For Active Directory Application Mode, you can back up data using the ADAM Writer. Then, you can perform a redirected restore of the data from the ADAM backup to the destination computer.

See [“About the Agent for Microsoft Active Directory”](#) on page 1311.

For more information, refer to your Microsoft documentation.

Table 5-3 How to install a new Windows Server domain controller into an existing domain by using a redirected restore

Step	Description
Step 1	Back up the System State data of an active Windows Server domain controller that is in the target domain. You should back up the data to some type of removable storage, such as a disk cartridge device or a tape. See “Backing up data” on page 158.
Step 2	Attach the storage that contains the System State data to the computer that you want to install into the destination domain. Note: It is recommended that you encrypt the storage. Use caution when transporting it to the location of the destination domain.
Step 3	Inventory and catalog the storage. See “Inventorying and cataloging a storage device” on page 563.
Step 4	Redirect the restore of the system state backup sets to a temporary location on a volume or directory on the destination computer. See “Restoring System State ” on page 247.
Step 5	Start the domain controller installation by doing the following in the order listed: ■ On the destination computer, click Start, and then click Run. ■ Type <code>dcpromo /adv</code>, and then click OK. ■ Click Next when the Active Directory Installation Wizard appears. ■ Select Additional domain controller for an existing domain, and then click Next. ■ Select From these restored backup files, enter the temporary location to which you redirected the restore of the System State data, and then click Next. ■ Complete the Active Directory Installation Wizard by following the prompts on the screen.
Step 6	Complete the domain controller installation.

Table 5-3 How to install a new Windows Server domain controller into an existing domain by using a redirected restore (*continued*)

Step	Description
Step 7	Restart the computer that has the new domain controller.
Step 8	Expire any remaining system state backup sets that you redirected to the temporary location. See “Changing the expiration date of backup sets on disk-based storage” on page 362.

Restoring Backup Exec Shadow Copy Components

The Backup Exec Shadow Copy Components file system uses Microsoft’s Volume Shadow Copy Service to protect third-party application and user data on Windows computers. You can restore the items in Backup Exec Shadow Copy Components individually or together.

The following items are contained in Backup Exec Shadow Copy Components:

- Backup Exec Deduplication Disk Storage
- Distributed File System Replication (DFSR)
- OSISoft PI Server data

When you restore domain-based Microsoft Distributed File System (DFS) namespaces, you must also restore the Active Directory to the same point-in-time. Domain-based DFS namespaces reference information that resides in the Active Directory. If you restore the namespaces without restoring Active Directory to the same point-in-time, that information is not restored and you may receive errors in the DFS Management Console.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

To restore Backup Exec Shadow Copy Components

- 1 On the **Backup and Restore** tab, right-click the server, and then click **Restore**.
- 2 Select **Shadow Copy Components**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“Restoring System State ”](#) on page 247.

Restoring utility partitions or Unified Extensible Firmware Interface system partitions

You can select utility partitions or Unified Extensible Firmware Interface (UEFI) partitions for restore. Utility partitions are small partitions that OEM vendors such as Dell, Hewlett-Packard, and IBM install on the disk. These partitions contain system diagnostic and configuration utilities. UEFI partitions are the small partitions that the operating system creates. The UEFI system partitions contain the critical system files, such as bootmgr and BOOT\BCD files.

Requirements for restoring utility partitions are as follows:

- You must recreate the utility partitions before you restore any data.
- You must have Administrator rights.
- You cannot redirect the restore of a utility partition to another computer.
- You can only restore the utility partitions that belong to the same vendor. For example, you cannot restore Dell utility partitions to a Compaq utility partition.
- The size of the utility partition to which you restore the data must be equal to or greater in size than the utility partition that was backed up.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

To restore utility partitions or UEFI system partitions

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore utility partitions or UEFI system partitions, and then click **Restore**.
- 2 Select one of the following, and then click **Next**:
 - **Utility partition**
 - **Unified Extensible Firmware Interface system partition**
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

About restoring encrypted data

Encrypted backup sets are identified in the restore selection list by an icon with a lock on it. When you select encrypted data to restore, Backup Exec automatically validates the encryption key for the data. If the encryption key that was used to back up the data is still in the Backup Exec Database, then Backup Exec selects that encryption key automatically. However, if the encryption key cannot be located, Backup Exec prompts you to provide the pass phrase for the encryption key that was used to back up the data. If you enter the correct pass phrase, Backup Exec recreates the key.

When you use a restricted encryption key to back up data, users other than the key owner must enter the pass phrase to restore data.

See [“Using encryption with Backup Exec”](#) on page 729.

See [“Encryption key management”](#) on page 733.

About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec

Backup Exec supports restoring NetWare SMS volume backups to non-SMS volumes. For example, the data that is backed up with Backup Exec for NetWare Servers or Novell’s SBackup can be restored to the Backup Exec server or to another network share.

You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

Canceling a restore job

Warning: Canceling a restore job while it is in progress may result in unusable data and can leave the disk in an unusable state.

To avoid canceling a restore job, you can redirect the restore to a noncritical destination. Then, copy the data to a final destination when the job completes successfully.

You should not cancel a System State restore job. Canceling a System State restore job can leave the system unusable.

To cancel a restore job

- 1 On the **Job Monitor**, right-click the job that you want to cancel.
- 2 Click **Yes** when you are prompted if you are sure that you want to cancel the job.

See [“Canceling an active job”](#) on page 265.

How Backup Exec catalogs work

While backing up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. When you select data to restore, Backup Exec uses the catalog information to find the restore selections and the storage devices on which they reside.

When a storage device is fully cataloged, you can do the following:

- View information on all the directories and files that are contained in each backup set.
- Search for files to restore.

Backup Exec catalogs each backup job. However, if the catalogs are truncated, only backup set information is listed. You cannot view files or file attributes. The amount of information in the catalog is determined by whether you choose to truncate the catalogs after a specific amount of time.

Catalogs reside on the Backup Exec server and on the storage device to which you sent the backup job.

To restore the data that was backed up by another installation of Backup Exec, you must first run a catalog operation on the storage device on the local Backup Exec server. The catalog for a backup job that was run on one installation of Backup Exec does not exist on another installation of Backup Exec.

When you enable Granular Recovery Technology (GRT) for Exchange, SharePoint, or virtual machine backups, a catalog operation runs immediately after the backup operation by default.

For Exchange and SharePoint agent-based backups, the full catalog operation runs immediately after all full backups. It runs once every 24 hours for all incremental backups and differential backups, even if you schedule more than one GRT-enabled job to run in the 24-hour period.

For Hyper-V and VMware backups, the full catalog operation runs immediately after all full, incremental, and differential backups by default. You can also schedule a full catalog operation.

For information about the best practices to manage catalogs in Backup Exec, refer to *Backup Exec Best Practices*.

See [“Configuring default options for catalogs”](#) on page 255.

See [“Cataloging a storage device”](#) on page 560.

See [“About cataloging tape or disk cartridge media that contains encrypted backup sets”](#) on page 520.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Configuring default options for catalogs

You can configure the default options for catalogs to specify the defaults that are best suited for your environment.

See [“How Backup Exec catalogs work”](#) on page 254.

To configure default options for catalogs

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Catalog**.
- 3 Edit any of the following options:

Request all media in the sequence for catalog operations

Catalogs the media in tape drives and disk cartridges by starting with the lowest known tape number in the tape family. For example, if you don't have the first tape, the catalog job starts with the second tape. If you uncheck this option, the catalog job begins on the tape that you specify.

If you uncheck **Request all media in the sequence for catalog operations**, then you cannot select the option **Use storage-based catalogs**.

This option is enabled by default.

Use storage-based catalogs

Lets Backup Exec read the catalog information from the storage device.

Storage-based catalogs allow fast cataloging of the backup sets that are not included in the Backup Exec server-based catalog. An example is when you want to catalog backup sets that another installation of Backup Exec creates.

Storage-based catalogs enable backup sets to be cataloged in minutes, rather than the hours that are required with traditional file-by-file cataloging methods.

To create a new catalog by having Backup Exec read each file block, clear this option. You should clear this option only if normal catalog methods are unsuccessful.

Note: It is recommended that you always attempt to use storage-based catalogs first before clearing this option.

If you uncheck **Request all media in the sequence for catalog operations**, then the option **Use storage-based catalogs** is unavailable.

This option is enabled by default.

The **Use storage-based catalogs** option must be enabled for the following restore scenarios:

- If you use the NDMP feature, this option must be enabled so that NAS server backup sets can be cataloged.
See ["About restoring and redirecting restore data for NDMP servers"](#) on page 1408.
- If you use the Advanced Disk-based Backup feature, this option must be enabled to use the true image restore capability.
See ["About true image restore for synthetic backups"](#) on page 1386.
- If you use Simplified Disaster Recovery, this option must be enabled to use the backup sets as part of an SDR recovery operation.
See ["Preparing computers for use with Simplified Disaster Recovery"](#) on page 881.
- If you back up a virtual machine to tape, this option must be enabled to perform a Granular Recovery Technology enabled restore with the backup sets.

See [“Granular Recovery Technology”](#) on page 739.

- If you back up Microsoft 365 tenants, this option must be enabled to ensure that the catalog job is successful.

See [“Backing up Microsoft 365 tenant data”](#) on page 433.

Truncate catalogs after

Retains only the header information and removes all file details and directory details after the specified amount of time. This option reduces the size of the catalogs considerably. After the catalogs have been truncated, the files and directories cannot be restored until you recatalog the storage.

See [“Cataloging a storage device”](#) on page 560.

The last access date is not reset when catalogs are truncated.

You can perform a full restore of backup sets from truncated catalogs.

This option does not apply to synthetic backup jobs.

This option is not enabled by default.

Catalog path

Designates a path on the volume for the catalog files. If the path does not exist, you are prompted to create the path.

The default path is C:<*Backup Exec install path*>\Backup Exec\Catalogs.

Note: It is recommended that you use Backup Exec Utility to change the catalog location. If you change the location using the **Catalog path** field, you must also manually copy the existing catalogs to the new location, and then restart the Backup Exec services.

See [“Moving the Backup Exec catalogs to a new directory”](#) on page 258.

4 Click **OK**.

Moving the Backup Exec catalogs to a new directory

It may be necessary to move the catalogs to a different location due to disk space limitations or other issues. Backup Exec can be configured to use a different directory for the catalog files.

The directory in which catalogs are stored from the Backup Exec server can also be changed. When changing the catalog directory from the Backup Exec server, all Backup Exec services must be stopped first and manually copy any existing catalogs to the new catalog directory.

To move the Backup Exec catalogs to a new directory

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Catalog**.
- 3 In the **Catalog** path field, browse to or type the new path in which to store catalogs.

Note: The directory must exist in Windows. Create the directory using Windows Explorer before you type it in the **Catalog path** field.

- 4 Click **OK**.
- 5 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Services**.
- 6 Click **Stop all services**.
- 7 Click **OK**.

After stopping all Backup Exec services, manually copy the catalog files from the old folder to the new one.

- 8 Using Windows Explorer, navigate to the directory in which the catalogs are currently stored.
- 9 Copy all of the items in the folder, including any subfolders and all files.
- 10 Using Windows Explorer, navigate to the directory to which you want to move the catalogs.
- 11 Paste all of the items in the folder.

After any catalog files have been copied to the new directory, restart all Backup Exec services.

- 12 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Services**.
- 13 Click **Start all services**.
- 14 Click **OK**.

Cataloging backup sets

Before you can restore or verify data, the data must be cataloged. While backing up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. However, you can manually catalog backup sets at any time.

To catalog backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to catalog.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To catalog a single backup set, right-click the backup set.
 - To catalog multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Catalog**.

Backup Exec runs the catalog operation.

See [“Backup sets ”](#) on page 359.

See [“How Backup Exec catalogs work”](#) on page 254.

Job management and monitoring

This chapter includes the following topics:

- [How to monitor and manage jobs in Backup Exec](#)
- [About the Job Monitor](#)
- [About the Job History](#)
- [Viewing the job log](#)
- [Error-handling rules for failed or canceled jobs](#)
- [List of job statuses in Backup Exec](#)
- [Setting job status and recovery options](#)
- [About Anomaly Detection](#)
- [Managing anomaly detection](#)

How to monitor and manage jobs in Backup Exec

Backup Exec provides many ways to monitor and manage your backup, restore, and storage operation jobs.

Job monitoring

You can monitor your active jobs and scheduled jobs from the **Job Monitor** tab or from the **Jobs** list on the **Backup and Restore** tab or the **Storage** tab. You can monitor the types of jobs that are active and scheduled, the state and status of the jobs, the schedule, and other details. The status of reports can be monitored from the **Reports** tab.

Table 6-1 Places where you can monitor jobs

Location where you can monitor jobs	More information
The Job Monitor tab	The Job Monitor provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The Job Monitor is divided into two panes. The Jobs pane displays the details of all of your active jobs and scheduled jobs. The Job Histories pane displays the details about all of your jobs that ran recently. By default, all recent jobs appear in the Job Histories pane. You can change the default by filtering the Job Histories view. Note: Report jobs do not appear on the Job Monitor. To monitor and manage report jobs, go to the Reports tab.
The Backup and Restore tab	When you double-click a server name on the Backup and Restore tab, details for that server become available. You can view the backup and restore jobs that are scheduled and active for the selected server. You can also see the history of backup and restore jobs that were processed on the selected server.
The Storage tab	When you double-click a storage device name on the Storage tab, details for that storage device become available. You can view the storage operation jobs that are scheduled and active for the selected storage device. You can also see the history of storage operation jobs that were processed for the selected storage device.
The Reports tab	On the Reports tab, you can see the scheduled reports in the See Upcoming report group, and you can see a list of reports that completed in the See Completed report group.

Job management

Backup Exec includes many features that enable you to manage backup jobs, restore jobs, and storage operation jobs.

You can manage jobs in the following ways:

- Edit scheduled jobs.
- Delete scheduled jobs.
- Cancel active jobs.
- Change the priority of scheduled jobs.
- Run a scheduled job immediately.
- Hold a job or the job queue.
- Run a test of the backup job.
- View the job activity details, such as job statistics and system information, for an active job.
- View the history of a job.
- Retry only failed resources.
- Run backup job with debugging enabled.
- Manage anomalies.

See [“Editing backup definitions”](#) on page 205.

See [“Viewing job activity details for active jobs”](#) on page 265.

See [“Deleting scheduled jobs”](#) on page 272.

See [“Canceling an active job”](#) on page 265.

See [“Changing the priority for a scheduled job”](#) on page 270.

See [“Running a scheduled job immediately”](#) on page 270.

See [“Holding jobs”](#) on page 266.

See [“Holding the job queue”](#) on page 268.

See [“Running a test run job manually”](#) on page 226.

About the Job Monitor

The **Job Monitor** provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The **Job Monitor** is divided into two panes. The **Jobs** pane displays the details of all of your active jobs and scheduled jobs. The **Job History** pane displays the details about all of your jobs that ran recently. By default, all recent jobs appear in the **Job History** pane. You can change the default by filtering the **Job History** view.

Note: Report jobs do not appear on the **Job Monitor**. To monitor and manage report jobs, go to the **Reports** tab.

You can manage active and scheduled jobs in the **Jobs** pane on the **Job Monitor** in the following ways:

- Edit scheduled jobs.
- Delete scheduled jobs.
- Cancel active jobs.
- Change the priority of scheduled jobs.
- Run a scheduled job immediately.
- Hold a job or the job queue.
- Run a test backup job.
- View the job activity details for an active job.
- View the history of a job.
- Retry only failed resources.
- Run backup job with debugging enabled.

In the **Job History** pane of the **Job Monitor**, you can manage completed jobs in the following ways:

- Run the job again immediately.
- View the job log.
- Delete the job.
- View the job history details.
- Enable error-handling rules.
- Duplicate job histories.
- Run a verify backup job.
- Edit the settings for a backup job.
- Retry only failed resources.
- Run backup job with debugging enabled.

From the **Job Monitor**, you can restore data from the backup sets of a completed backup job. Additionally, you can view the backup calendar.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 261.

See [“Running a test run job manually”](#) on page 226.

See [“Duplicating backup sets or a job history manually”](#) on page 222.

See [“Verifying backed up data manually”](#) on page 227.

See [“Editing backup definitions”](#) on page 205.

Viewing job activity details for active jobs

When a job is running, you can view details for the job, such as the percent complete, job rate, and byte count.

To view details for active jobs

- 1 Do one of the following:

To view job activity from the **Job Monitor** tab Select the **Job Monitor** tab.

To view job activity from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device on which the job is running.
- In the left pane, click **Jobs**.

- 2 Right-click the job, and then click **View Job Activity**.

Canceling an active job

You can cancel a job that is in progress. If the job is scheduled, it runs again at the next scheduled time.

It may take several minutes for a job to cancel. While Backup Exec processes the cancellation of a job, the Cancel Pending status appears in the **Job Status** column.

To cancel an active job

- 1 Do one of the following:

To cancel the job from the **Job Monitor** tab Select the **Job Monitor** tab.

To cancel the job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running.
- In the left pane, click **Jobs**.

- 2 Right-click the active job that you want to cancel, and then click **Cancel**.
- 3 Click **Yes** to confirm the cancellation of the job.

See [“List of job statuses in Backup Exec”](#) on page 289.

Holding jobs

Active and scheduled jobs can be placed on hold. When you place an active job on hold, the job continues to run until it is complete. However, the next scheduled occurrence of that job is placed on hold. When you place a job on hold, the status in the **Job Status** column changes to On Hold.

To hold jobs

- 1 Do one of the following:

To place a job on hold from the **Job Monitor** tab

Select the **Job Monitor** tab.

To place a job on hold from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running or is scheduled to run.
- In the left pane, click **Jobs**.

2 Do one of the following:

To hold a single job

Select the job from the list of jobs, and then in the **Jobs** group, click **Hold**. Then select **Hold Job**.

To hold all active jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Active Jobs**. You may need to click **Yes** to confirm that you want to hold all of the active jobs.

Note: This option is available only from the **Job Monitor**.

To hold all scheduled jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Scheduled Jobs**. You may need to click **Yes** to confirm that you want to hold all of the scheduled jobs.

Note: This option is available only from the **Job Monitor**.

See [“Removing the hold on jobs”](#) on page 267.

Removing the hold on jobs

You can remove the hold on a job at any time. When you remove the hold on a job, the status in the **Job Status** column changes to the job's original status, such as Active or Scheduled.

To remove the hold on jobs

1 Do one of the following:

To remove the hold on jobs from the **Job Monitor** tab

Select the **Job Monitor** tab.

To remove the hold from jobs from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is on hold.
- In the left pane, click **Jobs**.

2 Do one of the following:

To remove the hold on a single job

Select the job from the list of jobs, and then in the **Jobs** group, click **Hold**. Then select **Hold Job** to remove the check mark.

Note: If the job that you want to hold is part of a backup definition, you may need to double-click the job definition to view the job.

To remove the hold on all active jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Active Jobs** to remove the check mark. You may need to click **Yes** to confirm that you want to remove the hold on all of the active jobs.

Note: This option is available only from the **Job Monitor**.

To remove the hold on all scheduled jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Scheduled Jobs** to remove the check mark. You may need to click **Yes** to confirm that you want to remove the hold on all of the scheduled jobs.

Note: This option is available only from the **Job Monitor**.

See [“Holding jobs”](#) on page 266.

Holding the job queue

You can place the entire job queue on hold to make changes to your environment. The server is paused to place the job queue on hold. When the job queue is on hold, only active jobs continue to run unless you choose to cancel them. No other jobs can run until the job queue is taken off hold.

To place the job queue on hold

- 1 Do one of the following

To place the job queue on hold from the **Job Monitor** tab Select the **Job Monitor** tab.

To place the job queue on hold from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running or scheduled to run.
- In the left pane, click **Jobs**.

- 2 In the **Jobs** group, click **Hold**, and then click **Hold Job Queue**.
- 3 Click **Yes**.
- 4 If active jobs are running, select the active jobs that you want to cancel, and then click **OK**.

See [“Removing the hold on the job queue”](#) on page 269.

Removing the hold on the job queue

When you remove the hold on the job queue, the server is unpaused and jobs then run according to the schedule.

To remove the hold on the job queue

- 1 Do one of the following

To remove the hold on the job queue from the **Job Monitor** tab Select the **Job Monitor** tab.

To remove the hold on the job queue from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job queue is on hold.
- In the left pane, click **Jobs**.

- 2 In the **Jobs** group, click **Hold**, and then click **Hold Job Queue** to remove the check box.

See [“Holding the job queue”](#) on page 268.

Running a scheduled job immediately

You can run a scheduled job immediately. The job will also run on the next scheduled occurrence.

To run a scheduled job immediately

- 1 Do one of the following:

To run the job from the **Job Monitor** tab Select the **Job Monitor** tab.

To run the job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job that you want to run, and then click **Run Now**.

Note: If the scheduled job is part of a backup definition, you may need to double-click the job definition to view the scheduled job.

See [“Changing the priority for a scheduled job”](#) on page 270.

See [“Deleting scheduled jobs”](#) on page 272.

Changing the priority for a scheduled job

The priority determines the order that jobs run. If two jobs are scheduled to run at the same time, the priority you set determines which job runs first. The priority is changed for all occurrences of the scheduled job.

The priority of the job is displayed in the **Priority** column in the **Jobs** list.

This option is most useful if there are limited storage devices in your environment, but you want certain jobs to have priority access to the devices. A ready job that has a higher priority runs before a ready job that has a lower priority. A ready job that has a higher priority also runs before a ready job that has an earlier scheduled start time.

If multiple jobs are ready to run but must wait for a storage device to become available, then Backup Exec determines which jobs to run first. Backup Exec reviews the job priority and the scheduled start time of the job.

You can set the job priority when you do the following tasks:

- Select storage options when you create or edit a backup job
- View scheduled jobs in the **Job Monitor** tab
- View a server's scheduled jobs from the **Jobs** pane on the **Backup and Restore** tab

If you change the job priority in the storage options or on the **Jobs** pane, you can choose from the following levels of priority:

- **Highest**
- **High**
- **Medium**
- **Low**
- **Lowest**

If you change the job priority from the **Job Monitor**, you can only increase or decrease the job priority. You cannot specify a level.

To change the priority for a scheduled job

- 1 Do one of the following:

To change the job's priority from the **Job Monitor** tab Select the **Job Monitor** tab.

To change the job's priority from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job, and then click **Change Priority**.

Note: If the job is part of a backup definition, you may need to double-click the job definition to view the job.

- 3 Select the new priority.

Deleting scheduled jobs

Deleting a scheduled job removes all scheduled occurrences of the job. To delete only the occurrence of a scheduled job on a specific date, you can edit the schedule to remove that date.

Note: If a backup definition includes more than one type of job, then you cannot use the **Delete** option to delete an individual job from the definition. Instead, you must edit the job definition to remove the scheduled job.

To delete a scheduled job

- 1 Do one of the following:

To delete a scheduled job from the **Job Monitor** tab

Select the **Job Monitor** tab.

To delete a scheduled job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job, and then click **Delete**.

Note: If the job is part of a backup definition, you may need to double-click the job definition to view the job.

- 3 Click **Yes**.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 261.

Retrying only failed resources

You can retry a failed backup job to backup only the resources that failed during the immediate previous run of that backup job. This option is available only for a failed or cancelled backup job. If a backup job is successful, this option is grayed-out.

For example, if you have a VMware backup job that is backing up 5 virtual machines. During backup, only 3 virtual machines are successfully backed up and for 2 virtual machines the backup failed. You can refer to the job log to know

the reason for the failure, correct the issue, and then use the **Retry Only Failed Resources** option to run the job again. This time only the 2 virtual machines that failed during the immediate previous run of that backup job are backed up. The 3 virtual machines that were backed up successfully are not backed up again.

If a backup job is enabled for Simplified Disaster Recovery and failure occurs in one or more of the critical resources, then when you retry the failed backup job, all the critical resources are backed up again. Even the critical resources that may have backed up successfully in the immediate previous run of that backup job.

In a CAS-MMS environment, for jobs delegated from the CAS server, the option to run retry failed backup is only available from the CAS user interface. The option is disabled for the delegated jobs on the MMS user interface.

To retry only failed resources

- 1 Do one of the following:

To retry the failed resource from the **Job Monitor** tab. Select the **Job Monitor** tab.

To retry the failed resource from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running.
- In the left pane, click **Jobs**.

- 2 Do one of the following:

- Right-click the failed job that you want to retry, and then click **Retry Only Failed Resources**.
- On the **Job Monitor** tab, select the failed job that you want to retry, and then click **Retry Only Failed Resources**.

A dialog box is displayed stating that the job will run immediately and also run as per the defined schedule.

- 3 Click **Yes**.

The job starts and only the resource that failed during the immediate previous run of the job is backed up.

Run backup job with debugging enabled

You can run a backup and restore job with automatic debugging enabled. When you call technical support with an issue in your backup or restore job, technical

support asks you to gather the debug logs. With the help of the debug logs, support can then make changes to the jobs that have an issue. After the job is complete, you have all debug logs that you need to share with technical support.

To run a job with debugging enabled

- 1 Do one of the following:

To run the job with debugging enabled from the **Job Monitor** tab

Select the **Job Monitor** tab.

To run the job with debugging enabled from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running.
- In the left pane, click **Jobs**.

- 2 Do one of the following:

- Right-click the job , and then click **Run with Debugging Enabled**.
- On the **Job Monitor** tab, select the job, and then click **Run with Debugging Enabled**.

If the job does not have a schedule defined, a dialog box is displayed stating that the job will run immediately.

If the job has a schedule defined, a dialog box is displayed and you must select if you want the job to run immediately or as per the schedule. The debugging enabled job runs only the first time that you enable the option. It does not run on the subsequent schedules.

- 3 Click **OK**.

The job is completed and the debug logs are generated. If the backup or restore job has an issue, you can send these debug logs to technical support.

About the Job History

The **Job History** displays a list of completed and failed backup, restore, and storage operation jobs. The **Job History** appears on the lower pane of the **Job Monitor** tab. It also appears when you select a server on the **Backup and Restore** tab, and when you select a storage device on the **Storage** tab.

From the **Job History**, you can do any of the following:

- View the job log.

- Delete a job.
 - Rerun a job.
 - Duplicate the data from a completed backup job.
 - Verify a backup job.
 - Enable error-handling rules for a failed job.
- See [“Running a job from the Job History”](#) on page 276.
- See [“Viewing the job log ”](#) on page 279.
- See [“Deleting a job from the Job History”](#) on page 275.
- See [“Enabling an error-handling rule for a failed job”](#) on page 287.
- See [“Duplicating backup sets or a job history manually”](#) on page 222.
- See [“Verifying backed up data manually”](#) on page 227.

Viewing the history of a job

The job history shows statistics for all occurrences of a job.

To view the history of a job

- 1 Do one of the following:

To view the history of a job from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To view the history of a job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then click **View Job History**.

Deleting a job from the Job History

You can delete a job from the **Job History**, or have Backup Exec automatically delete the job history using database maintenance.

If you delete a job, it is removed from the computer and cannot be recovered.

To delete a job from the Job History

- 1 Do one of the following:

To delete a job from the Job History on the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To delete a job from the Job History on the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

2 Right-click the job that you want to delete, and then click **Delete**.

You can select multiple jobs by selecting a job, and then pressing the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This lets you perform tasks such as Delete on more than one job at a time, as long as the jobs are of similar type.

You can delete up to 2500 jobs from the **Job History**. If you attempt to delete more than 2500 jobs, you are prompted to continue with the deletion.

3 Click **Yes**.

See [“About the Job History”](#) on page 274.

See [“Configuring database maintenance and security”](#) on page 701.

Running a job from the Job History

After a job runs, the job moves to the **Job History**. You can run a completed job again from the **Job History**.

To run a job from the Job History

1 Do one of the following:

To run a job from the **Job History** on the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To run a job from the **Job History** on the **Backup and Restore** tab or the **Storage** tab

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

2 Right-click the job that you want to run, and then click **Run Now**.

See [“About the Job History”](#) on page 274.

Retrying only failed resources from the Job History

You can retry a job that failed from the job history for only the resources that failed during the immediate previous run of that backup job. This option is available only for a failed or cancelled backup job. If a backup job is successful, this option is grayed-out.

For example, if you have a VMware backup job that is backing up 5 virtual machines. During backup, only 3 virtual machines are successfully backed up and for 2 virtual machines the backup failed. You can refer to the job log to know the reason for the failure, correct the issue, and then use the **Retry Only Failed Resources** option to run the job again. This time only the 2 virtual machines that failed during the immediate previous run of that backup job are backed up.. The 3 virtual machines that were backed up successfully are not backed up again.

If a backup job is enabled for Simplified Disaster Recovery and failure occurs in one or more of the critical resources, then when you retry the failed backup job, all the critical resources are backed up again. Even the critical resources that may have backed up successfully in the immediate previous run of that backup job.

In a CAS-MMS environment, for jobs delegated from the CAS server, the option to run retry failed backup is only available from the CAS user interface. The option is disabled for the delegated jobs on the MMS user interface.

To retry only failed resources from the Job History

1 Do one of the following:

To retry the failed resource from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To retry the failed resource from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

2 Do one of the following:

- Right-click the failed job that you want to retry, and then click **Retry Only Failed Resources**.
- On the **Job Monitor** tab, select the failed job that you want to retry, and then click **Retry Only Failed Resources**.

A dialog box is displayed stating that the job will run immediately and also run as per the defined schedule.

3 Click **Yes**.

The job starts and only the resource that failed during the immediate previous run of the job is backed up.

Run backup job with debugging enabled from the Job History

You can run a backup and restore job with automatic debugging enabled. When you call technical support with an issue in your backup or restore job, technical support asks you to gather the debug logs. With the help of the debug logs, support can then make changes to the jobs that have an issue. After the job is complete, you have all debug logs that you need to share with technical support.

To run a job with debugging enabled

1 Do one of the following:

To run the job with debugging enabled from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To run the job with debugging enabled from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running.
- In the left pane, click **Job History**.

2 Do one of the following:

- Right-click the job , and then click **Run with Debugging Enabled**.
- On the **Job Monitor** tab, select the job, and then click **Run with Debugging Enabled**.

If the job does not have a schedule defined, a dialog box is displayed stating that the job will run immediately.

If the job has a schedule defined, a dialog box is displayed and you must select if you want the job to run immediately or as per the schedule. The debugging enabled job runs only the first time that you enable the option. It does not run on the subsequent schedules.

3 Click **OK**.

The job is completed and the debug logs are generated. If the backup or restore job has an issue, you can send these debug logs to technical support.

Viewing the job log

You can view detailed job-related properties for each job that has been processed. You can save a copy of the job log to a location of your choice or you can print the job log.

To view the job log

- 1 Do one of the following:

To view the job log from the Job Monitor tab	On the Job Monitor tab, locate the job in the Job Histories pane.

To view the job log from the Backup and Restore tab or the Storage tab	Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.
- 3 Do any of the following:
 - Click **Find** to search for a particular word, phrase, or file name.
 - Click **Save As** to save a copy of the job log to the location of your choice.
 - Click **Print** to print the job log.

Finding text in the job log

You can search for specific text in the job log. Backup Exec searches only the sections that are expanded. To search the entire job log, you should select the **Expand All** option.

To find text in the job log

- 1 Do one of the following:

To find text in the job log from the Job Monitor tab	On the Job Monitor tab, in the Job Histories pane, locate the job for which you want to search the job log.

To find text in the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then click **View Job Log**.
- 3 Click **Find**.
- 4 In the **Find** field, type the text that you want to find.
- 5 Select any of the following additional options to facilitate your search:

Match whole word only

Check this check box if you want Backup Exec to search for the whole word you typed. If you do not select this option, Backup Exec finds the text that includes part of the word. For example, if you search for the word "file" and do not select this option, Backup Exec finds all occurrences of "file", "files", "filed", and any other words that contain "file". If you do select this option, Backup Exec finds only the occurrences of "file".

Match case

Check this check box if you want Backup Exec to search for words using the exact capitalization that you typed. For example, if you search for the word "File" and select this option, Backup Exec finds all occurrences of "File", but does not find any occurrences of "file".

Highlight all matches

Check this check box if you want Backup Exec to highlight the text that matches the search criteria. The option is selected by default.

- 6 Click **Next** to find the next occurrence of the text.

See [“Viewing the job log”](#) on page 279.

Printing the job log

If your Backup Exec server is connected to a printer, you can print the job log for a completed job.

To print the job log

- 1 Do one of the following:

To access the job log from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To access the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.
- 3 Click **Print**.
- 4 Select the printer that you want to use, and then click **Print**.

See [“Saving the job log”](#) on page 281.

Saving the job log

Backup Exec provides the ability to save the job log to a location on your hard drive or network. In addition, you can select the format in which to save the file; as a complete webpage, a web archive, an HTML-only webpage, or a text file.

To save the job log

- 1 Do one of the following:

To access the job log from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To access the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.

- 3 Click **Save As**.
 - 4 Select the location where you want to save the job log.
- See [“Printing the job log”](#) on page 281.

Linking from the job log to the Arctera Technical Support website

Errors that are reported in the job log have a unique code, called a Unique Message Identifier (UMI). These codes contain hyperlinks that you can click to go to the Arctera Technical Support website. From the website, you can access technical notes and troubleshooting tips that are related to a specific message. UMI codes establish unique message codes across all Arctera products.

Some alerts also contain a UMI. For example, if a Warning alert appears when a job fails, the alert includes the UMI code.

You can create or enable an error-handling rule for errors. These rules let you set options to retry or stop a job when the error occurs.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

To link from the job log to the Arctera Technical Support website

- 1 Do one of the following:

To link to the job log from the **Job Monitor** tab

Select the **Job Monitor** tab.

To link to the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click a job, and then select **View Job Log**.
- 3 Scroll to the **Job Completion Status** section.
- 4 Click the UMI code, which appears as a blue hyperlink.

How to use job logs with vertical applications

The Backup Exec Administration Console provides a view of the job logs in HTML format. If necessary, you can convert the job logs to a text format for use with vertical applications.

To convert a job log file to a text format, load the Backup Exec Management Command Line Interface, and then type the following at a command prompt:

Get-BEJobLog "pathname\job log filename"

For example, to display the job log C:<Backup Exec install path>\Backup Exec\Data\bex00001.xml in text format to the command prompt, you would type:

```
Get-BEJobLog "C:<Backup Exec install path>\Backup Exec\Data\bex00001.xml"
```

To redirect the job log to a file, you would type one of the following:

```
Get-BEJobLog "C:<Backup Exec install path>\Backup Exec\Data\bex00001.xml"  
> bex00001.txt
```

See [“Viewing the job log”](#) on page 279.

Configuring default job log options

You can configure default options for job logs that specify the amount of detail you want to include in the completed job log. For the jobs that produce large job logs, you may want to reduce the amount of detail in the job log. The size of the job log increases proportionally to the level of detail that is configured for the job log.

To configure default job log options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Job Logs**.
- 3 Select the appropriate options.

Summary information only

Select this option to include the following information in the job log:

- Job name
- Job type
- Job log name
- Backup Exec server name
- Storage device
- Starting date and time
- Errors encountered
- Ending date and time
- Completion statistics

This option also includes the names of files that were skipped, the name of the media set, the backup type and results of the verify operation if one was performed.

Summary information and directories processed

Select this option to include summary information and a list of all processed subdirectories in the job log.

Summary information, directories, and files processed

Select this option to include summary information, processed subdirectories, and a list of all the file names that were processed in the job log.

Summary information, directories, files, and file details

Select this option to include summary information, processed subdirectories, a list of all the file names and their attributes in the job log. This option increases the job log size significantly.

There may be an issue when you open the job log from the Backup Exec user interface because of Windows limitations. It is recommended that you use this option only during troubleshooting.

Prefix for the job log file name	Enter a prefix to add to the job logs that are processed. The default prefix is BEX. The job log file name consists of Prefix_ServerName_Count. Prefix is the label that you enter in this field, ServerName is the name of the Backup Exec server that ran the job, and Count is the number of job logs that this job has produced
Attach job logs as HTML	Select this option to attach the job logs in an HTML format when an email notification is sent.
Attach job logs as text	Select this option to attach the job logs in a text format when an email notification is sent.
Job log path	Shows the current location of the job log. To change the path you can use BE Utility.

See [“Viewing the job log ”](#) on page 279.

Error-handling rules for failed or canceled jobs

You can enable default rules or create custom rules to set retry options and final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you either place the job on hold until you can fix the error, or reschedule the job for its next scheduled service.

Each default error-handling rule applies to one category of errors, such as Network Errors or Security Errors. Default error-handling rules are disabled by default, so you must edit a rule and enable the rules that you want to use. You cannot delete default error-handling rules, add specific error codes to a category, or add new error categories. Before the error-handling rules will apply, the final error code must be in an error category that is associated with a rule, and the rule must be enabled.

To apply an error-handling rule for a specific error code that is in an error category, you can create a custom error-handling rule. You can select up to 28 error codes in an error category that a custom error-handling rule can apply to. You can also add an error code to an existing custom rule.

A custom error-handling rule named "Recovered Jobs" is created when Backup Exec is installed and is enabled by default. This rule applies retry options and a final job disposition to jobs that fail and that are not scheduled to run again.

See ["Creating a custom error-handling rule"](#) on page 286.

If both a custom error-handling rule and a default error-handling rule apply to a failed job, the settings in the custom rule are applied to the job.

Note: Error-handling rules are not applicable for duplicate jobs.

Creating a custom error-handling rule

You can create custom rules to set retry options and final job disposition for failed or canceled jobs.

See ["Error-handling rules for failed or canceled jobs"](#) on page 285.

To create a custom error-handling rule

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Click **New**.
- 3 Check **Enable error-handling rule**.
- 4 In the **Name** field, type a name for this rule.
- 5 In the **Error category** list, select the category of error that you want this rule to apply to.
- 6 Under **Available errors**, check the check box of the error code that you want this rule to apply to. You can select up to 28 error codes.
- 7 If you want Backup Exec to retry the job after it fails, check **Retry job**, and then enter the number of times you want to retry the job and how often you want to retry the job.

The maximum number of retries is 99. The maximum retry interval is 1440 minutes.

- 8 Under **Final job disposition**, select the way you want to handle the job after the maximum number of retries.

If you choose the option to place the job on hold until the error condition is cleared manually, you must manually remove the hold from the job after you manually clear the error condition.

- 9 Click **OK**.

See ["Custom error-handling rule for recovered jobs"](#) on page 288.

Enabling or disabling error-handling rules

Follow these steps to enable or disable specific error-handling rules.

To enable or disable error-handling rules

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Select the rule that you want to enable or disable, and then click **Edit**.
- 3 Do one of the following:
 - To enable the rule, check **Enable error-handling rule**.
 - To disable the rule, clear the **Enable error-handling rule** check box
- 4 Click **OK**.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

Deleting a custom error-handling rule

A custom error-handling rule can be deleted at any time. A default error-handling rule cannot be deleted.

To delete a custom error-handling rule

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Select the custom rule that you want to delete, and then click **Delete**.
- 3 Click **Yes** to confirm that you want to delete the rule.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

Enabling an error-handling rule for a failed job

You can create custom rules to set retry options and final job disposition for failed jobs.

To enable an error-handling rule for a failed job

- 1 Do one of the following:
- To enable an error-handling rule from the **Job Monitor** tab

Select the **Job Monitor** tab.

To enable an error-handling rule from the **Backup and Restore** tab or the **Storage** tab

On the **Backup and Restore** tab or the **Storage** tab, in the left pane, select **Job History**.
- 2 Right-click the failed job, and then select **Error Handling**.
- 3 Check the **Enable error-handling rule** check box.
- 4 Complete the remaining options for this rule.

Custom error-handling rule for recovered jobs

Backup Exec includes a custom error-handling rule called "Recovered Jobs" to recover the jobs that failed with specific errors. This rule is created when Backup Exec is installed and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. The final job disposition is to place the job on hold until you have manually cleared the error condition.

The following table describes the error codes that are selected by default for the Recovered Jobs custom error-handling rule.

Table 6-2 Error codes for recovered jobs custom error-handling rule

Error code	Description
0xE00081D9 E_JOB_ENGINE_DEAD	The displayed error message is: The Backup Exec job engine system service is not responding.
0xE0008820 E_JOB_LOCAL RECOVERNORMAL	The displayed error message is: The local job has been recovered. No user action is required.
0xE000881F E_JOB_REMOTE RECOVERNORMAL	The displayed error message is: The remote job has been recovered. No user action is required.

Table 6-2 Error codes for recovered jobs custom error-handling rule
(continued)

Error code	Description
0xE0008821 E_JOB_STARTUP RECOVERY	The displayed error message is: Job was recovered as a result of Backup Exec RPC service starting. No user action is required.

Note: If the Central Admin Server feature is installed, additional error codes are selected.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

List of job statuses in Backup Exec

Backup Exec displays various job statuses for active, completed, and scheduled jobs.

See [the section called “Active job statuses”](#) on page 289.

See [the section called “Completed job statuses”](#) on page 291.

See [the section called “Scheduled job statuses”](#) on page 292.

Active job statuses

The following statuses may appear for the jobs that are active:

Table 6-3 Active job statuses

Status	Description
Running	The operation is underway.
Queued	The job has been initiated, but Backup Exec is actively looking for a suitable drive or media.
Cancel Pending	Backup Exec cannot process the Cancel request immediately. This status is displayed until the job is actually canceled. The job is then displayed in the job history with a status of Canceled.
Loading Media	The media is being loaded and positioned on the target device.

Table 6-3 Active job statuses (*continued*)

Status	Description
Pre-processing	This status can indicate any or all of the following: ■ Backup Exec is calculating the amount of data that will be backed up, if the Display progress indicators for backup jobs option is enabled in the Preferences section of Backup Exec settings. See “Changing the default preferences” on page 698. ■ Backup Exec is waiting for a pre-command or a post-command to complete. ■ Backup Exec is retrieving the set maps and is positioning the tape to the append point location for an append job.
Snapshot processing	Backup Exec is processing a snapshot operation.
Device Paused	The device that the job was sent to is paused. See “Pausing and unpausing a storage device” on page 564.
Server Paused	The Backup Exec server is paused. See “Pausing or resuming a managed Backup Exec server” on page 1370.
Stalled	The Backup Exec services have become unresponsive. See “Setting job status and recovery options” on page 295.
Media Request	You must insert media for the job to continue.
Communication Stalled	Communications between the managed Backup Exec server and the central administration server have not occurred within the configured time threshold. See “Enabling or disabling communications between the managed Backup Exec server and the central administration server” on page 1358.
No Communication	No communication about jobs is being received at the central administration server from the managed Backup Exec server. The configured time threshold has been reached. See “Enabling or disabling communications between the managed Backup Exec server and the central administration server” on page 1358.
Consistency check	Backup Exec is running a consistency check of the databases before backup.

Table 6-3 Active job statuses (*continued*)

Status	Description
Updating Catalogs	Backup Exec is updating the catalog information.

Completed job statuses

The following statuses may appear for the jobs that were completed:

Table 6-4 Job completion status

Status	Description
Successful	The job completed without errors.
Completed with exceptions	The job completed, but some files were in use, skipped, or corrupted.
Resumed	The status is the same as the failed over status, however the Enable checkpoint restart option was selected.
Canceled	The administrator terminated the operation as it was running.
Canceled, timed out	The Cancel the job if it is still running x hours after its scheduled start feature was enabled and the job was not completed within the specified time frame.

Table 6-4 Job completion status (*continued*)

Status	Description
Failed	The operation took place, but one or more significant errors occurred. The job log should indicate what caused the errors so that you can decide if you want to run the job again. For example, if a job failure occurred due to a lost connection during job processing, you could choose to resubmit the job when the connection is restored. If a drive loses power during a backup operation, you should restart the backup job using a different tape. You can restore the data that was written to the tape up to the point of the power loss, but you should not reuse the tape for subsequent backup operations. A failed job has an error message in the Errors section of the job log with a link to the Arctera Technical Support website. A job may fail for the following reasons: ■ The storage device that was selected for the job was not available when the job ran. ■ The logon account that was used in the backup job is incorrect. Verify that the logon account information is valid for the resource being backed up. ■ A problem occurred with the storage device when the job ran. ■ The computer being backed up was shut down before or during the backup job.
Recovered	The job was active when the status of the managed Backup Exec server was changed from Communication Stalled to No Communication. The custom error-handling rule for Recovered Jobs was applied to the job.
Missed	The job did not run during the scheduled time window. The job is rescheduled to run based on the time window that you configured.

Scheduled job statuses

The following statuses may appear for the jobs that are scheduled:

Table 6-5 Scheduled job statuses

Status	Description
Invalid Schedule	The scheduled job will not run because of a scheduling issue. See “Setting default schedule options for rule-based jobs and run now jobs” on page 692.
Not in time window	The job was ready to be sent for processing, but the time window for the job closed. See “Setting default schedule options for rule-based jobs and run now jobs” on page 692.
On Hold	The job has been placed on hold.
Queued	A temporary state that displays when Backup Exec is applying an error-handling rule that is enabled to retry the job. See “Custom error-handling rule for recovered jobs” on page 288.

Table 6-5 Scheduled job statuses (*continued*)

Status	Description
Ready	The job is ready to run, but cannot for one of the following reasons: ■ Internal error. No devices are available, but the cause is unknown. ■ Invalid job. The job type is unknown; there may be an internal error or the database is corrupted. ■ Invalid target. This type of storage device no longer exists. ■ Backup Exec server not available. ■ No license for option name. A license must be purchased and installed on the Backup Exec server. ■ No Backup Exec servers are available. ■ No Backup Exec servers are available in Backup Exec server pool. ■ Specified destination storage device pool is empty. ■ Specified destination device is not in Backup Exec server pool. ■ Specified destination device not on local Backup Exec server. ■ Specified destination storage device pool on local Backup Exec server is empty. ■ The destination storage device cannot be a storage pool. ■ The destination storage device cannot be a Backup Exec server. ■ Another job is running in the system that is blocking execution of this job. This job will run after the other job completes. ■ Invalid input. ■ Incompatible Resumes. ■ No server license available. ■ No multi-server license available. ■ No Windows license. ■ No Windows server. ■ Need local Backup Exec server. ■ Local server is not a Backup Exec server. ■ No idle storage devices are available. ■ No eligible storage devices within the storage pool are available. ■ Blocked by an active, linked duplicate backup sets job.

Table 6-5 Scheduled job statuses (*continued*)

Status	Description
Scheduled	The job is scheduled to run in the future. The scheduled jobs that are linked to another job, such as a job to duplicate backup sets, will not display a scheduled job status.
Server Paused	The job is ready, but the Backup Exec server has been paused. No jobs are dispatched while the Backup Exec server is paused. See “Pausing or resuming a managed Backup Exec server” on page 1370.
To Be Scheduled	A state that the scheduled job transitions through as it is being sent for processing.
Rule Blocked	The job cannot run because one or more of the settings in the backup definition cannot be satisfied. The Rule Blocked status may appear for any of the following reasons: ■ A job cannot run until another job has completed. Example: If you added a duplicate stage to a backup definition and the source backup job has not yet completed, the duplicate job is blocked until the source backup job completes. Example: If a job definition includes both a full backup and an incremental backup, the full backup must run first. If you try to run the associated incremental backup job before the full backup job has completed, the incremental backup job is blocked until the full backup completes. ■ A server for a job cannot be changed until the linked jobs complete. ■ A server for a job cannot be changed until another job completes. ■ A job cannot run because multiple jobs are scheduled to run at the same time.

Setting job status and recovery options

If the Backup Exec services become unresponsive or jobs no longer run, you can set the threshold at which Backup Exec changes the status of active jobs to stalled.

You can also set the threshold at which Backup Exec fails the jobs that were stalled, and then recovers them.

By lowering the amount of time before Backup Exec reaches the threshold for changing a job's status to stalled, you can receive an earlier notification that jobs have stalled. A shorter time between the stalled and recovered thresholds also allows Backup Exec to fail and then recover the stalled jobs earlier. However, setting the thresholds too low may force a job to be recovered when it is not necessary.

Backup Exec recovers the jobs by using the custom error-handling rule named Recovered Jobs. This custom error-handling rule is created and enabled when Backup Exec is installed, and specifies that stalled, failed, and recovered jobs are retried two times, with an interval of five minutes between the retries.

Jobs that are stalled and then failed and recovered by Backup Exec because of unresponsive Backup Exec services are displayed differently in Backup Exec than the jobs that fail because of errors in normal daily activities. The stalled/failed/recovered jobs are not indicated in red text in the job history as other failed jobs are. Instead, these jobs are displayed in gray text with a job status of **Recovered**.

In the job history, the error category is listed as Job Errors. The job history indicates the type of internal communication error that occurred and that the job was recovered. Based on the type of error that occurred, a log file may or may not be associated with the recovered job.

To set job status and recovery options

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Job Status and Recovery**.

3 Set the thresholds for stalled and recovered jobs.

Stalled

Enter the amount of time you want to wait before Backup Exec changes an unresponsive job's status to Stalled.

Recovered

Enter the amount of time you want to wait before Backup Exec fails jobs that stalled and then recovers them. A custom error-handling rule named Recovered Jobs is applied to recovered jobs. If this rule is disabled, then any other error-handling rules that have been enabled will apply to the recovered jobs. If no error-handling rules apply to the job, then the job fails.

4 Click **OK**.

See [“List of job statuses in Backup Exec”](#) on page 289.

See [“Custom error-handling rule for recovered jobs”](#) on page 288.

About Anomaly Detection

The Anomaly Detection feature monitors backup job parameters and then detects and reports anomalies for the backup job. No anomalies are detected until a minimum of 30 runs are completed for a backup job.

The feature also allows you ignore patterns that help to reduce detection of records that are not anomalies. In such a scenario, when the job runs again, that pattern is not detected as an anomaly.

The following backup job parameter, attributes, features, and so on are verified during anomaly detection:

- Backup image size
- Backup Item count
- Deduplication Ratio
- Backup Time
- Data transferred

From the Backup Exec user interface, you can see the anomalies for the last 30 days. You can filter anomalies for the past 7 days, 15 days, or 30 days.

After the anomalies are detected, a score is calculated, based on which the anomalies are categorized as low severity, medium severity and high severity. You can manage the anomalies by confirming that a record is an anomaly or confirming that the record is not an anomaly.

See [“Configuring the Home tab”](#) on page 120.

See [“Managing anomaly detection”](#) on page 298.

You can disable this feature from the Backup Exec settings. After it is disabled, Backup Exec stops detecting any anomalies.

See [“Changing the default preferences”](#) on page 698.

You can also generate the Anomaly Detection Summary report to view the details of the anomalies for the backup jobs.

See [“Anomaly Detection Summary report”](#) on page 809.

Managing anomaly detection

You can view the list of reported anomalies to confirm a record as an anomaly or to confirm that it is not an anomaly. By default, the list of anomalies is displayed for the last 7 days. To view anomalies for more than 7 days, you can apply filters.

Each page in the **Anomaly Detection** window displays 100 records. Any selection or action that you perform on the anomalies is limited to that specific page.

To manage anomaly detection

- 1 Do any of the following:
 - In the **System Health** group, under **Anomaly Detection**, click **Manage Anomalies**.
 - On the **Job Monitor** tab, click **Anomaly Detection > Anomaly Information**.

On the **Anomaly Detection** window, the list of anomalies is displayed.

- 2 For an anomaly, you can do the following:

Job Name	Displays name of the backup job.
Resource Name	Displays name of the resource that is backed up. The resources can be backup sets or jobs (Microsoft 365).
Detection Time	Date and time when the anomaly is detected.

Severity	Displays the severity of the anomaly. The anomaly score is calculated based on the backup job parameters, attributes, or features. Based on score the severity of the anomaly is categorized. ■ Score: 0-10 - Low ■ Score: 10-15 - Medium ■ Score: > 15 - High
Anomaly Status	Displays the status of the anomaly. ■ False Positive ■ Anomaly Confirmed
Reason	Reason given for the False Positive anomaly status.
Image Size	Displays size of the backup image and usual size range of the image.
Backup Item Count	Displays number of items that are backed up and usual range of the number of items that are backed up.
Deduplication Ratio	Displays ratio of the amount of data before deduplication to the amount of data after deduplication. Usual range of the deduplication ratio is also displayed.
Backup Time	Displays backup time in seconds and usual range of the backup time.
Data Transferred	Displays total size of data transferred over the network and usual size range of the data transferred. For deduplication storage devices the size can be less than the actual image size. For other storage devices the size is same as the backup image size.
Score	Displays the anomaly score.

- Report as False Positive** This can be done in two ways:
- Select the checkbox corresponding to the record and click **Report as False Positive** to confirm that the record is not an anomaly.
 - Select multiple records and click **Report as False Positive** to confirm that the records are not anomalies.

During multiple selections, the option is in a disabled state if at least one record is confirmed as an anomaly. Remove the record from the selection and confirm that the selected records are not anomalies.

You must also enter a reason why any of the selected records are not anomalies.

- Confirm as anomaly** This can be done in two ways:
- Select the checkbox corresponding to the record and click **Confirm as anomaly** to confirm that the record is an anomaly.
 - Select multiple records and click **Confirm as anomaly** to confirm that the records are anomalies.

During multiple selections, the option is in a disabled state if at least one record is confirmed as not an anomaly. Remove the record from the selection and confirm that the selected records are anomalies.

View Job Log Click to view the job log and job history.

Refresh Click to update the anomaly list. Any actions that you performed are removed.

Collapse All Click to collapse all expanded records.

- 3 Click **Apply Filter** to filter and view the anomalies.
By default, anomalies for the last 7 days are displayed.

- 4 On the **Manage Anomaly Detection** window, select the filters and click **OK**.

Job Name	Select Enable this filter to filter records based on the job name.
Resource Name	Select Enable this filter to filter records based on the resource name.
Detection Time	Select Enable this filter to filter records based on the time period. You can filter the records for a selected number of days or select a date range. Ensure that the date range does not exceed 30 days.
Severity	Select Enable this filter to filter records based on the severity of the anomaly.

- 5 Click **OK**.

Alerts and notifications

This chapter includes the following topics:

- [Alerts and notifications in Backup Exec](#)
- [Where to find alerts in Backup Exec](#)
- [Enabling active alerts and alert history to display on the Home tab](#)
- [Viewing the alert history for a server or a storage device](#)
- [Deleting an alert from alert history](#)
- [Copying alert text to a document or email](#)
- [Filtering alerts](#)
- [Viewing the job log from an alert](#)
- [Responding to active alerts](#)
- [Clearing all informational alerts manually](#)
- [Setting up notification for alerts](#)
- [Configuring email or text message notification for alerts](#)
- [Adding recipients for notification](#)
- [Adding a recipient group for alert notifications](#)
- [Disabling email or text message alert notification for a recipient](#)
- [Editing recipient notification properties](#)
- [Deleting recipients](#)
- [Configuring alert categories](#)

- [Assigning recipients to receive notifications for specific alert categories](#)
- [Sending a notification when a job completes](#)
- [Disabling notifications for a specific alert category](#)
- [Configuring default alert settings](#)
- [Enabling or disabling pop-up alerts](#)
- [SNMP traps for Backup Exec alerts](#)

Alerts and notifications in Backup Exec

An alert is any event or condition in Backup Exec that is important enough to display a message or require a response from you. Backup Exec includes many alert categories and four alert types. Alert categories are the events or the conditions that cause alerts. Alert categories encompass many circumstances or problems that affect the system, jobs, media, or storage sources. Each alert category can include one or more events that generate an alert. For example, a Job Failed error may occur for many reasons. The alert severity levels can help you to determine which alerts need immediate attention and which alerts require a response.

The following alert severity levels are used in Backup Exec:

Table 7-1 Alert severity

Item	Description
Attention required	Indicates the issues that require a response before the job or operation can continue.
Error	Indicates the issues that affect job processing or the integrity of your backup. These alerts cannot be disabled and cannot be configured to be cleared automatically. You must respond to them manually.
Warning	Indicates the conditions that may or may not cause jobs to fail. You should monitor the conditions and take actions to resolve them.
Informational	Provides status messages for the conditions that you might want to know about.

By default, most alerts are enabled, which means that they appear in the **Active Alerts** pane when they occur. You can disable warning alerts and informational

alerts by editing alert category properties. However, error alerts and attention required alerts cannot be disabled. You can filter the alerts so that only specific alerts appear.

See [“Configuring alert categories”](#) on page 319.

From the **Home** tab, you can view all active alerts or filter the alerts to view only specific alert severities or only the alerts that occurred on certain dates. On the **Backup and Restore** tab, when you double-click a server, you can see the active alerts that are specific to that server. Similarly, on the **Storage** tab, when you double-click a type of storage, you can see the active alerts that are specific to that storage device.

Alerts remain in the **Active Alerts** pane until they receive a response. You can respond to an alert manually or you can configure Backup Exec to respond to some alerts automatically after a specified length of time. Depending on the alert severity, a response might not be required, such as with informational alerts. After you respond to an alert, Backup Exec moves it to the alert history. Alert history is available on the **Home** tab, the **Backup and Restore** tab, and the **Storage** tab. In addition, an Alert History report is available from the **Reports** tab.

See [“Alert History report”](#) on page 808.

See [“Responding to active alerts”](#) on page 310.

You can configure notifications to inform recipients when alerts occur. For example, you can notify a backup administrator by email or cell phone text message when a critical alert occurs.

See [“Setting up notification for alerts”](#) on page 311.

To assist with hardware troubleshooting, Backup Exec displays alerts for SCSI event ID 9 (storage timeout), ID 11 (controller error), and ID 15 (storage not ready).

Where to find alerts in Backup Exec

You can find alerts in multiple locations on the Backup Exec Administration Console.

Table 7-2 Where alerts appear on the Backup Exec Administration Console

Location of alerts	Details
Home tab	To see a list of active alerts on the Home tab, ensure that the Active Alerts check box is checked in the System Health group. To see a list of alerts in alert history, ensure that the Alert History check box is checked in the System Health group. From the Active Alerts widget on the Home page, you can respond to alerts, view the job log, clear all informational alerts, or view the details of specific alerts.
Backup and Restore tab	On the Backup and Restore tab, active alerts are listed for each server on the list of servers. Also, when you double-click a server, you can select Active Alerts in the left pane to display all of the active alerts for that server. When you select Active Alerts in the left pane, the Show Alert History option also becomes available in the Alerts group at the top of the dialog box. Additionally, you can click the alert icon next to a server name to view the alerts for that server.
Storage tab	On the Storage tab, active alerts are listed for each storage device on the list of storage devices. Also, when you double-click a storage device, you can select Active Alerts in the left pane to display all of the active alerts for that device. When you select Active Alerts in the left pane, the Show Alert History option also becomes available in the Alerts group at the top of the dialog box. Additionally, you can click the alert icon next to a storage device name to view the alerts for that device.

Table 7-2

Where alerts appear on the Backup Exec Administration Console

(continued)

Location of alerts	Details
Backup Exec status bar	The Backup Exec status bar shows the number of active alerts for each type of alert. Double-click in the Alerts section of the Backup Exec status bar to see a list of all the active alerts and the alert history, along with details of those alerts.
Windows status bar	If you have minimized Backup Exec, the Backup Exec icon on the Windows status bar flashes when a Backup Exec alert is triggered. When you maximize Backup Exec, the pop-up alerts appear.
Pop-up window	Alert pop-up windows appear when an alert is triggered. The color of the alert pop-up corresponds to the type of alert; red for error alerts, yellow for warning alerts, purple for attention-required alerts, and blue for informational alerts. If more than three alerts are available, the pop-up lists the number of alerts that are pending and the color of the pop-up indicates the most severe of the alert types. For example, if three informational alerts and two error alerts have occurred, the pop-up alert indicates that five alerts have arrived. Additionally, the pop-up color will be red to indicate that error alerts are part of the group of new alerts. You can respond to attention required alert types directly from the pop-up.

Enabling active alerts and alert history to display on the Home tab

The **Active Alerts** pane appears on the **Home** tab by default. If it does not appear, follow these steps to show the alert details. Optionally, you can also enable a history of all alerts for the server to appear on the **Home** tab.

To view active alerts and alert history on the Home tab

- ◆ On the **Home** tab, in the **System Health** group, do any of the following:

- Check the **Active Alerts** check box to see a list of active alerts.
- Check the **Alert History** check box to see a list of all alerts that occurred on the server.

See “[Active alerts properties](#) ” on page 307.

Active alerts properties

Properties for active alerts can be viewed on the **Home** tab or in the details for a backup job, a restore job, or a type of storage.

See “[Enabling active alerts and alert history to display on the Home tab](#)” on page 306.

Table 7-3 Properties for active alerts

Item	Description
Severity	Indicates the severity of the alert. The severity helps you determine how quickly you want to respond. The following alert severity levels may appear: ■ Error ■ Warning ■ Information ■ Attention Required
Category	Indicates the condition that caused the alert. Categories include Database Maintenance, General Information, Device Error, or Job Failed.
Message	Indicates the text of the error message.
Date and Time	Shows the date and time when the alert was received.
Job Name	Indicates the name of the job that triggered the alert. This column is blank if a job triggered the alert, such as for general information alerts.
Storage	Shows the name of the storage device on which the alert occurred.
Server	Shows the name of the server on which the alert occurred.

Table 7-3 Properties for active alerts *(continued)*

Item	Description
Source	Indicates the cause of the alert. Alerts can originate from one of the following sources: ■ System ■ Job ■ Storage ■ Media

Viewing the alert history for a server or a storage device

After you respond to an alert, Backup Exec moves it to the alert history.

To view the alert history for a server or a storage device

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the item for which you want to view the alert history.
- 2 In the left pane, click **Active Alerts**.
- 3 In the **Alerts** group, select **Show Alert History**.

Deleting an alert from alert history

Follow these steps to delete an alert from the alert history.

To delete an alert from alert history

- 1 Access the alert history from any of the following locations:
 - In the Backup Exec status bar, double-click in the **Alerts** area.
 - On the **Backup and Restore** tab, double-click a server. In the left pane, select **Active Alerts**, and then click **Show Alert History**.
 - On the **Storage** tab, double-click a storage device. In the left pane, select **Active Alerts**, and then click **Show Alert History**.
- 2 On the **Alerts** dialog box, select the **Alert History** tab.
- 3 Right-click the alert that you want to remove from the alert history, and then click **Delete**.
- 4 Click **Yes** to confirm that you want to delete the alert.

Copying alert text to a document or email

Backup Exec gives you the ability to copy alert information to a document, such as Notepad or Microsoft Word, or to an email. When you copy alert text to a word-processing application or to an email, Backup Exec formats the alert text in a table. Alert text that is copied to Notepad displays in plain text.

To copy the text of an alert to a document or email

- 1 Access the list of alerts from any of the following locations:
 - In the Backup Exec status bar, double-click in the **Alerts** area.
 - On the **Backup and Restore** tab, double-click a server. In the left pane, select **Active Alerts**.
 - On the **Storage** tab, double-click a storage device. In the left pane, select **Active Alerts**.
- 2 On the **Alerts** dialog box, select either the **Active Alerts** tab or the **Alert History** tab, depending on whether the alert you want to copy is active or in the alert history.
- 3 Right-click the alert that you want to copy, and then click **Copy**.

Note: You can also select an alert, and then press Ctrl + C as a shortcut to copy the alert text. In addition, you can copy multiple alerts by using Shift + Click or Ctrl + Click.

- 4 Open the document or email to which you want to copy the alert text, right-click in the document or email, and then select **Paste**.

Filtering alerts

You can filter the alerts that appear in the **Active Alerts** pane on the **Home** tab. Filters are useful when you have many alerts and you want to only view specific alert types. Alerts can be filtered by severity, time, and source. For example, you can choose to view only the error alerts that occurred during the last 12 hours for jobs.

To filter alerts

- 1 On the **Home** tab, locate the **Active Alerts** pane.
If the **Active Alerts** pane does not appear, you must enable the alert details.
See [“Enabling active alerts and alert history to display on the Home tab”](#) on page 306.
- 2 Use any combination of the following options to filter the alerts list:
 - In the **Source** field, select the source of the alerts that you want to view.
 - In the **Time** field, select the time frame for which you want to view alerts.
 - In the **Severity** field, select the severity levels of the alerts that you want to view, such as **Error** or **Warning**.

Viewing the job log from an alert

The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status for completed jobs. You can access the job log from the alerts that were generated for jobs.

To view the job log from an alert

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click the alert for which you want to view the job log, and then select **View Job Log**.
- 3 Do any of the following:
 - To search for a specific word or phrase, click **Find**. Type the text you want to find, and then click **Next**.
Be sure to expand all sections of the job log. The Find feature searches only the expanded sections of the job log.
 - To print the job log, click **Print**. To print the log, you must have a printer attached to your system and configured.
 - To save the job log as an .html file or a .txt file, click **Save As** and then select the file name, file location, and file type.

Responding to active alerts

You can respond to active alerts and continue or cancel the operation, depending on the alert condition. By default, Backup Exec displays all enabled alerts, and

all alerts that require a response. If you have set filters, only those alerts that are selected appear in addition to any alerts that require a response.

If you click **Close** on the alert response dialog box, the dialog box closes, but the alert remains active. To clear the alert, you must select a response such as **OK**, **Yes**, **No**, or **Cancel**. You can configure automatic responses for some alert categories.

See [“Configuring alert categories”](#) on page 319.

Some alerts provide a Unique Message Identifier (UMI) code. This code is a hyperlink to the Arctera Technical Support website. You can access the technical notes that are related to the alert.

To respond to an active alert

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click the alert that you want to respond to, and then click **Respond** or **Respond OK**.
- 3 Click a response for the alert, such as **Respond OK** or **Respond**.

Clearing all informational alerts manually

You can configure individual alert categories to be cleared automatically after a certain period of time. Informational alerts may be generated often, so you may want to clear all informational alerts manually before the system moves them automatically.

To clear all informational alerts manually

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click an informational alert, and then select **Clear All Informational Alerts**.

See [“Configuring alert categories”](#) on page 319.

Setting up notification for alerts

You can configure Backup Exec to notify recipients when alerts occur. Setting up notifications for alerts is a three-step process.

Table 7-4 How to set up notification for alerts

Step	Action
Step 1	Configure the method you want to use to notify the recipient. The notification methods are text message or email. See “Configuring email or text message notification for alerts” on page 312.
Step 2	Add the contact information for the people or groups that you want to receive notifications. See “Adding recipients for notification” on page 314. See “Adding a recipient group for alert notifications” on page 316.
Step 3	Assign each recipient to the receive notifications for specific alert categories. See “Assigning recipients to receive notifications for specific alert categories” on page 321. See “Configuring alert categories” on page 319. See “Sending a notification when a job completes” on page 321.

Configuring email or text message notification for alerts

You can set up Backup Exec to send email or text messages to specified recipients when an alert occurs. Email notification requires an email account to be used as the sender. For example, you might want to use an email account for the backup administrator or the IT administrator. To configure email notifications, enter the name of the sender's mail server, the port number that the server uses, and the sender's name and email address. You can also set up Backup Exec to authenticate the emails that are sent for alerts.

Table 7-5 Types of notification

Notification type	Description
Email notification	Backup Exec uses SMTP for email notifications and supports authentication, Transport Layer Security (TLS), and Secure Sockets Layer (SSL). Notification email messages can be sent to Microsoft Outlook and web-based email applications, such as Gmail, Yahoo mail, or Microsoft Office 365 (O365).

Table 7-5 Types of notification (continued)

Notification type	Description
Text message notification	For a text message notification, Backup Exec attempts to format the message to contain fewer than 144 characters to meet text messaging protocol restrictions. By limiting a notification to fewer than 144 characters, the notification is more likely to be sent in a single text message instead of broken up into multiple messages. However, the text messaging service provider determines how the notifications are delivered. Text message notifications are sent in the following formats: ■ Job-related notification: Backup Exec: <Server Name> : <Job Name> : <Status> ■ Alert-related notification: Backup Exec: <Server Name> : <Alert Type>

After the sender's email information has been entered, then information about recipients can be set up.

Note: An SMTP-compliant email system, such as a POP3 mail server, is required for email notifications.

Note: After you configure email or text message notification, you cannot remove the configuration to disable notifications. However, you can disable notification for individual recipients.

To configure Backup Exec to send email or text notification for alerts

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Email and Text Notification**.
- 3 Set up the sender's email and/or text messaging information.

If you want to send alert notifications by email Do the following:

- Under **Email configuration**, enter the name of the mail server, the port number that the mail server uses, and the sender's name and email address.
- If you want to authenticate the email information that you entered, check **Enable email authentication**, and then enter the user name and password for the sender's email account.

If you want to send alert notifications by text message In **Text message service provider address**, enter the fully-qualified domain name of the sender's text messaging service provider.

Example: If a company called "MyPhone" provides text messaging services, then enter "MyPhone.com" in the **Text message service provider address** field. You can override this default address for the individuals who do not use this provider.

Note: Text message notification is sent as SMTP mail to an email address that is provided by a text messaging service provider. To enable notification by text message, you must enter the information about the sender's email account in the **Email configuration** section in addition to the fully-qualified domain name of the default text messaging service provider.

4 Click **OK**.

You can now add information about the recipients who need to receive alert notifications.

See [“Adding recipients for notification”](#) on page 314.

Adding recipients for notification

Individuals or groups can be set up to receive notifications when alerts occur in Backup Exec. When you set up an individual recipient, you indicate whether the

person wants to receive notifications by email, text message, or both. A group recipient contains the individual recipients that you select. Each individual within a group receives notifications by the method that is indicated for the individual; email, text message, or both.

Note: Information about the notification sender must be configured before recipients can be configured.

See [“Configuring email or text message notification for alerts”](#) on page 312.

To add recipients for email or text message notification

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, click **Add a recipient**.
- 4 In the **Name** field, type the name of the recipient.
- 5 Select the method of notification for this recipient:

To send emails to this recipient

Check **Send notifications by email**, and then enter the person's email address.

To limit the number of emails that are sent within a specific amount of time, check **Send no more than x emails within x minutes/hours**, and then enter the maximum number of emails to send in a specific time period.

To send text messages to this recipient

Check **Send notifications by text message**, and then enter the person's cell phone number and text message service provider's address.

Note: If the recipient is located in a different country, you must include the exit code of the country from which the message is generated and the country code for the country in which the recipient is located. For example, the exit code for the United States is 011. The country code for Italy is 39. To send a message from the United States to a recipient that is located in Italy, enter **011 39** and the recipient's phone number.

The phone number can include spaces and the following characters:

- Opening and closing quotes
- Period
- Plus sign
- Dash
- Opening and closing parentheses
- Forward slash

To limit the number of text messages that are sent within a specific amount of time, check **Send no more than x text messages within x minutes/hours**, and then enter the maximum number of text messages to send in a specific time period.

6 Click **OK**.

You can now assign the recipients to the alert categories for which they should receive notifications.

See [“Assigning recipients to receive notifications for specific alert categories”](#) on page 321.

Adding a recipient group for alert notifications

Groups are configured by adding recipients as group members. A group contains one or more recipients and each recipient receives the notification message. A group can only include individuals. A group cannot contain other groups.

To add a recipient group for alert notifications

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, click **Add a group**.
- 4 In the **Name** field, type a unique name for this notification group.
- 5 To add members to the group, select recipients from the **All recipients** list, and then click **Add** to move them to the **Selected recipients** list.

To remove members from the group, select recipients from the **Selected recipients** list, and then click **Remove** to move them to the **All recipients** list.

- 6 When you have completed the group, click **OK**.

The group is added to the list of recipients on the **Manage Recipients** dialog box.

See [“Adding recipients for notification”](#) on page 314.

Removing a recipient from a group

When you remove a recipient from a group, the recipient no longer receives the notifications that the group is configured to receive. The recipient continues to receive notifications for which the recipient is configured to receive as an individual recipient.

To remove a recipient from a group

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, double-click the group that contains the recipient.
- 4 Under **Selected recipients**, select the recipient that you want to remove, and then click **Remove**.

Disabling email or text message alert notification for a recipient

If a person no longer wants to receive alert notifications, you can disable notifications for that person.

Note: Instead of disabling notifications completely, you can also change the alert categories for which a recipient receives notifications.

See [“Disabling notifications for a specific alert category”](#) on page 322.

To disable email or text message alert notification for a recipient

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the name of the recipient, and then click **Edit**.
- 4 Do any of the following:
 - To disable email notifications, clear the **Send notifications by email** check box.
 - To disable text message notifications, clear the **Send notifications by text message** check box.
- 5 Click **OK**.

Editing recipient notification properties

You can edit the recipient notification properties at any time and change the recipient information, such as an email address or cell phone number. For a group, you can add recipients to the group or remove recipients from the group.

To edit the recipient notification properties

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the recipient that you want to edit.
- 4 Click **Edit**.
- 5 Edit the properties for the selected recipient.
- 6 Click **OK**.

Deleting recipients

You can delete the recipients that do not want to receive notification messages. The recipient is permanently removed upon deletion. Alternatively, you can disable notification for recipients.

See [“Disabling notifications for a specific alert category”](#) on page 322.

See [“Disabling email or text message alert notification for a recipient”](#) on page 317.

To delete a recipient

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the recipient that you want to delete.
- 4 Click **Delete**.
- 5 Click **Yes** to confirm that you want to delete this recipient.
- 6 Click **OK**.

Configuring alert categories

Alert categories are the events or the conditions that cause alerts. Alert categories encompass many circumstances or problems that affect the system, jobs, media, or storage sources. Each alert category can include one or more events that generate an alert. For example, a Job Failed error may occur for many reasons. The alert types can help you to determine which alerts need immediate attention and which alerts require a response. You can set up alert categories to enable or disable alerts and to determine what actions should take place when an alert occurs.

Most alerts are enabled by default, however the following alert categories are initially disabled:

- Backup job contains no data
- Job Start
- Job Success

You can disable alert categories if they are informational or warning alerts. You cannot disable alert categories if they are error or attention required alerts.

Each time you change the alert configuration, it is recorded in the audit log. You can view the audit log at any time to view the changes that were made to the alert category.

To configure alert category properties

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Under **Alert category**, select the alert that you want to configure.

4 Under **Category Properties**, select the appropriate options.

Enable alerts for this category	Enables or disables the alert. You cannot disable error and attention required alerts.
Include the job log with email notifications	Sends the job log to the recipient that is configured for notification. If you select this option, be sure to select at least one recipient from the Send notification to the following recipients area at the bottom of the dialog box.
Record event in the Windows Event Log	Enters the alert into the Windows Event Viewer. The Windows Event log displays all the property information for the alert. If a link does not appear in the Windows Event log you can search the Technical Support website for information about the Event ID.
Send SNMP notifications	Indicates whether SNMP notifications are enabled or cleared for the alert. SNMP must be installed to use this option.
Automatically clear after X hours/minutes	Lets you enter the number of minutes, hours, or days you want the alert to remain active before it is cleared. Note: Error alerts cannot be cleared automatically, so this option is disabled for error alerts.
Respond with	Indicates the response that you want Backup Exec to send automatically when the alert is cleared. This option is available only for the Media Overwrite and Media Insert alert categories and only when the Automatically clear after x days/hours/minutes option is selected. The choices are Cancel , No , Yes , or OK .

Send notification to the following recipients

Lets you select the name of a recipient to notify when this type of alert occurs. You must have recipients configured to use this option.

If the recipient to which you want to send notifications is not in the list, click **Manage Recipients** to add the recipient.

- 5 Repeat steps 2 - 4 to configure additional alert categories.
- 6 Click **OK** to save the properties that you selected.

Assigning recipients to receive notifications for specific alert categories

After you have set up notification recipients, you should determine the alert categories for which they should receive notifications. For example, some recipients may only want to receive notifications about job failures and tape errors while other recipients may want to receive notifications for all error alert categories.

To assign recipients to receive notifications for specific alert categories

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Select an alert category from the list.
- 4 If the category is not enabled, click **Enable alerts for this category**.
- 5 Under **Send notification to the following recipients**, check the check box next to the name of every recipient that needs to receive notifications for the selected alert category.
- 6 Click **OK**.

Sending a notification when a job completes

You can assign recipients to be notified when a job completes. Recipients must be set up before you can set up notification.

To send a notification when a job completes

- 1 Create a new job or edit an existing job.
- 2 On the **Backup Options** dialog box, in the left pane, click **Notification**.

- 3 Select the check box for each recipient that you want to notify when each type of job completes.
- 4 To send the job log with the notification to an email address, check **Include job log in email notifications**.
- 5 You can continue selecting other options from the **Options** dialog box or click **OK**.

Notification options for jobs

When you set up or edit a job, you can select recipients to receive notification when the job completes.

See [“Sending a notification when a job completes”](#) on page 321.

Table 7-6 Notification options for jobs

Item	Description
Recipient name	Shows the names of the individual and group recipients.
Recipient type	Indicates Recipient for an individual recipient or Group for a group recipient.
Include job log in email notifications	Enables Backup Exec to include a copy of the job log with the notification. This option applies only to email recipients. The maximum attachment size, in kilobytes, can be configured in the following registry key: HKLM\Software\Veritas\Backup Exec for Windows\Backup Exec\Server\Max Notification Attachment Size The attachment size can also be determined by the settings on your mail server.
Manage Recipients	Lets you add, edit, or delete recipients.
Properties	Lets you view or change the properties of a selected recipient.

Disabling notifications for a specific alert category

When a recipient no longer needs to receive notifications for an alert category, you can stop the notification.

To disable notifications for a specific alert category

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Under **Alert category**, select the category for which a recipient no longer needs to receive notification.
- 4 Under **Send notifications to the following recipients**, clear the check box next to the recipient for whom you want to stop notification.
- 5 Click **OK**.

See [“Disabling email or text message alert notification for a recipient”](#) on page 317.

Configuring default alert settings

The default alert settings let you enable or disable the pop-up alerts for the four types of alerts and also to determine how long those pop-up alerts remain on the screen. If you disable the pop-up alerts for a particular alert type, that alert type still appears in the list of active alerts in other parts of the Backup Exec Administration Console, unless you have filtered that same alert type from the list of active alerts.

To configure default alert settings

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Backup Exec Settings**, and then select **Alerts**.
- 3 If you want to receive a reminder alert to renew your license contracts, check **Send an alert reminder to renew your license contracts on**, and then enter a date.
- 4 In the **License usage threshold warning for sending alert** field, enter the percentage of license usage that you want to use as a threshold to receive an alert for renewal.

By default, the threshold is set to 85%.

Based on the threshold that you enter, the **Healthy** or **Warning** icon is displayed on the **Home** tab. This option is enabled when Capacity licenses, Instance-based licenses, or Microsoft 365 licenses are installed.

- 5 In the **Display pop-up alerts for the following alert types** group box, check the check boxes for the types of alerts that you want to see in a pop-up alert. Clear the check boxes for the types of alerts that you do not want to see in a pop-up alert.

Informational

Informational alerts provide status messages for the conditions that you might want to know about. They do not require a response and are not critical. The pop-up informational alerts have a blue color.

Error

Error alerts indicate the issues that affect job processing or the integrity of your backup. You must respond to them manually. The pop-up error alerts have a red color.

Warning

Warning alerts indicate the conditions that may or may not cause jobs to fail. You should monitor the conditions and take actions to resolve them. The pop-up warning alerts have a yellow color.

Attention required

Attention required alerts indicate the issues that require a response before a job or an operation can continue. You can respond to this type of alert directly from the pop-up. The pop-up attention-required alerts have a purple color.

- 6 In the **Number of seconds to display pop-up alerts** field, enter the amount of time you want pop-up alerts to remain on the screen.
- 7 Click **OK**.

Enabling or disabling pop-up alerts

By default, Backup Exec displays informational, error, warning, and attention required alert types in pop-up alert windows when an error occurs. You can change the default settings so that you see only the types of alerts that are important to you.

To enable or disable pop-up alerts

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Backup Exec Settings**, and then select **Alerts**.
- 3 In the **Display pop-up alerts for the following alert types** group box, check the check boxes for the types of alerts that you want to see in a pop-up alert. Clear the check boxes for the types of alerts that you do not want to see in a pop-up alert.
- 4 Click **OK**.

SNMP traps for Backup Exec alerts

SNMP (Simple Network Management Protocol) is a method by which a network can be monitored from a central location. SNMP-enabled network applications like Backup Exec report to an SNMP console (a management workstation). The console receives messages (traps) from Backup Exec regarding status and error conditions. An MIB is available in the WINNT\SNMP\language directory on the Backup Exec installation media that you can load into your SNMP console.

The Object Identifier prefix for Veritas is:

1.3.6.1.4.1.1302

Backup Exec SNMP traps (messages) have unique object IDs and may include up to four strings.

The following SNMP trap types are supported:

Table 7-7 SNMP traps

Trap Type	Object ID	String 1	String 2	String 3	String 4
Product Start	1302.3.1.1.9.1	Backup Exec: Application initializing	machine name	product, version, revision	
Product Stop	1302.3.1.1.9.2	Backup Exec: Application terminating	machine name	product, version, revision	
Job Canceled	1302.3.1.2.8.2	Backup Exec: Job canceled by Operator	machine name	job name	local or remote Operator name
Job Failed	1302.3.1.2.8.1	Backup Exec: Job failed	machine name	job name	detail message

Table 7-7 SNMP traps (continued)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Storage device requires human intervention	1302.3.2.5.3.3	Backup Exec: Storage device requires attention	machine name	job name	detail message
Robotic library requires human intervention	1302.3.2.4.3.3	Backup Exec: robotic library device requires attention	machine name	job name	detail message
Simplified Disaster Recovery Message	1302.3.1.4.2.1.2	SDR full backup success	machine name	job name	detail message
Backup Exec system error	1302.3.1.1.9.3	The application has encountered an error	machine name	job name	detail message
Backup Exec general information	1302.3.1.1.9.4	Information on normal events	machine name	job name	detail message
Job Success	1302.3.1.2.8.3	The job succeeded	machine name	job name	detail message
Job Success with exceptions	1302.3.1.2.8.4	The job succeeded, but there was a problem	machine name	job name	detail message
Job Started	1302.3.1.2.8.5	The job has started	machine name	job name	detail message
Job Completed with no data	1302.3.1.2.8.6	The job succeeded, but there was no data	machine name	job name	detail message
Job Warning	1302.3.1.2.8.7	The job has a warning	machine name	job name	detail message
PVL Device Error	1302.3.1.5.1.1.1	The device has encountered an error	machine name	job name	detail message
PVL Device Warning	1302.3.1.5.1.1.2	The device has encountered a warning	machine name	job name	detail message

Table 7-7 SNMP traps (continued)

Trap Type	Object ID	String 1	String 2	String 3	String 4
PVL Device Information	1302.3.1.5.1.1.3	Normal device information	machine name	job name	detail message
PVL Device Intervention	1302.3.1.5.1.1.4	Device requires attention	machine name	job name	detail message
PVL Media Error	1302.3.1.5.2.1.1	There is an error with the media	machine name	job name	detail message
PVL Media Warning	1302.3.1.5.2.1.2	There may be a problem with the media	machine name	job name	detail message
PVL Media Information	1302.3.1.5.2.1.3	Normal media information	machine name	job name	detail message
PVL Media Intervention	1302.3.1.5.2.1.4	Media requires attention	machine name	job name	detail message
Catalog Error	1302.3.1.5.3.1.1	There is an error with the catalog	machine name	job name	detail message
Tape Alert Error	1302.3.1.5.4.1.1	There is a TapeAlert error	machine name	job name	detail message
Tape Alert Warning	1302.3.1.5.4.1.2	There is a TapeAlert warning	machine name	job name	detail message
Tape Alert Information	1302.3.1.5.4.1.3	Normal TapeAlert information	machine name	job name	detail message
Database Maintenance Error	1302.3.1.5.5.1.1	There is a database maintenance error	machine name	job name	detail message
Database Maintenance Information	1302.3.1.5.5.1.2	Normal database maintenance information	machine name	job name	detail message
Install Update Warning	1302.3.1.5.7.1.1	There is an install warning	machine name	job name	detail message
Install Update Information	1302.3.1.5.7.1.2	Normal Install information	machine name	job name	detail message

See [“Installing and configuring the SNMP system service”](#) on page 328.

Installing and configuring the SNMP system service

To receive Backup Exec traps at the SNMP console, you must configure the SNMP system service with the SNMP console's IP address.

SNMP starts automatically after installation. You must be logged on as an administrator or a member of the Administrators group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

To install the SNMP system service and configure it to send traps to the SNMP console

- 1 From the Windows Control Panel, select **Add/Remove Programs**.
- 2 Click **Add/Remove Windows Components**.
- 3 In Add/Remove Windows Components, select **Management and Monitoring Tools**, and then click **Details**.

When selecting the component, do not select or clear its check box.

- 4 Select **Simple Network Management Protocol**, and then click **OK**.
- 5 Click **Next**.

Installing the Windows Management Instrumentation performance counter provider

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes performance counter and SNMP providers that can be manually installed and used with WMI.

To install the WMI performance counter provider

- 1 Insert the Backup Exec Installation media.
- 2 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\backupexecperfmon.mof
```

Installing the Windows Management Instrumentation provider for SNMP

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes

performance counter and SNMP providers that can be manually installed and used with WMI.

To use the WMI SNMP provider you must set up SNMP notification.

To install the WMI SNMP provider

- 1 Before you install the SNMP provider that is included with Backup Exec, you must have the Microsoft SNMP provider installed on your system.

For more information, refer to your Microsoft documentation.

- 2 Insert the Backup Exec Installation media.
- 3 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\snmp\eng\bkuexecmib.mof
```

Uninstalling the Windows Management Instrumentation performance counter provider

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI performance counter provider

- ◆ At the command line, type:

```
mofcomp <CD Drive  
Letter>:\winnt\wmi\deletebackupexecperfmon.mof
```

Uninstalling the Windows Management Instrumentation provider for SNMP

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI SNMP provider

- ◆ At the command line, type:

```
Smi2smir /d Backup_Exec_MIB
```

Disk-based and network-based storage

This chapter includes the following topics:

- [Features and types of disk-based storage and network-based storage](#)
- [Storage trending statuses for disk storage and virtual disks](#)
- [Setting low disk space thresholds on disk-based storage](#)
- [Configuring disk storage](#)
- [How to restore data from a reattached or reinserted disk-based storage device](#)
- [Configuring disk cartridge storage](#)
- [How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage](#)
- [Backup sets](#)

Features and types of disk-based storage and network-based storage

Features of disk-based storage include the following:

- Automatic discovery of locally accessible disk volumes.
- Disk space monitoring. Alerts are sent when the disk space thresholds that you set are reached.
- Storage trending analysis that provides predictions of low disk space for disk storage and virtual disks.

- Data lifecycle management, which automatically deletes expired backup sets and reclaims the disk space for use by new backup sets.

Disk-based storage includes the following types of storage:

Table 8-1 Types of disk-based storage

Types of disk-based storage	Description
Disk storage	Disk storage is a location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data. See “Configuring disk storage” on page 334.
Disk cartridge devices	Disk cartridges are a type of storage that usually remains attached to the Backup Exec server while you remove the media, such as RDX. If you are not sure if the storage has removable media, you can open the Computer folder on your Windows computer. The devices that contain removable media are listed. See “Configuring disk cartridge storage” on page 345.
Deduplication disk storage	Deduplication disk storage is a disk-based backup folder that is located on the Backup Exec server and which provides integrated deduplication. You must install the Backup Exec Deduplication feature to use this data-reduction strategy that optimizes storage and network bandwidth. See “About the Deduplication feature” on page 969.

Network-based storage includes NDMP servers, OpenStorage devices, cloud storage devices, and the Remote Media Agent for Linux.

Table 8-2 Types of network storage

Type of storage	Description
NDMP servers	NDMP servers are network-attached storage (NAS) that supports the Network Data Management Protocol (NDMP) to allow the use of devices that are attached to the servers. See “Features of the NDMP feature” on page 1399.
OpenStorage devices	OpenStorage devices are network-attached storage that supports Arctera's OpenStorage technology. See “Configuring an OpenStorage device” on page 419.
Cloud storage devices	Cloud storage devices are the storage devices that are configured on the cloud hosted by the cloud storage service provider.
Remote Media Agent for Linux	The Remote Media Agent for Linux lets you back up data from remote computers to the storage devices that are directly attached to a Linux server. You can also back up to a simulated tape library on a Linux server. See “About the Remote Media Agent for Linux ” on page 1456.

For information about the best practices to manage disk-based storage in Backup Exec, refer to *Backup Exec Best Practices*.

See [“Storage trending statuses for disk storage and virtual disks”](#) on page 332.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

See [“Backup sets ”](#) on page 359.

Storage trending statuses for disk storage and virtual disks

Backup Exec gathers disk usage information for disk storage and virtual disks. Backup Exec then performs statistical analysis of used disk space and free disk space. The analysis provides an estimate of how many days remain before the disk storage or virtual disk is full.

Alerts provide information about whether the current disk space resources are sufficient, and can help you plan when to increase disk space.

Table 8-3 Storage trending statuses

Storage trending status	Description
Remaining storage: x days	An estimate of the remaining number of days of storage space, based on the current usage of disk space.
History of used space is still being gathered	This status may appear for any of the following reasons: ■ The disk storage device has not been configured long enough to get a statistical estimate. Note: After you create disk storage, Backup Exec may take approximately one month to gather enough information to provide a storage estimate. ■ This storage may be on a managed Backup Exec server that is currently in a rolling upgrade.
Current storage is sufficient	The environment contains enough disk space to meet storage requirements for the next 30 days.
No estimate due to an inconclusive history of used space	A storage trend cannot be obtained. Unusual increases or decreases in the amount of free disk space in the last 30 days can cause this status.
Not enough statistical information is available	Backup Exec has not collected enough sample data for statistical analysis. Note: After you create disk storage, Backup Exec may take approximately one month to gather enough information to provide a storage estimate.

See [“Configuring disk storage”](#) on page 334.

Setting low disk space thresholds on disk-based storage

You can change the default values for three levels of low disk space conditions. When the storage device's used capacity reaches this threshold, Backup Exec

sends an alert, and the data lifecycle management feature immediately searches the device for expired backup sets that it can delete.

Data lifecycle management also runs on a disk cartridge if the cartridge reaches capacity during a backup job. The backup may not need to span to another cartridge if data lifecycle management deletes enough expired backup sets. If a job does span and you insert a new disk cartridge that is full, data lifecycle management deletes expired backup sets on the new cartridge.

To set low disk space thresholds on disk-based storage

- 1 On the **Storage** tab, double-click the storage on which you want to change the low disk space thresholds.
- 2 In the left pane, click **Properties**.
- 3 Change the value on any or all of the following properties:
 - **Low disk space.**
The first low disk space threshold at which you want Backup Exec to send an alert. The default value is 25%.
 - **Low disk space - Warning.**
The second low disk space threshold at which you want Backup Exec to send an alert. The default value is 15%. This threshold must be less than the Low disk space threshold.
 - **Low disk space - Critical.**
The third low disk space threshold at which you want Backup Exec to send an alert. The default value is 5%. This threshold must be less than the Warning threshold.
- 4 Click **Apply**.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Configuring disk storage

Disk storage is a location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data. You specify how long you want to keep the data that you back up to disk storage when you create a backup job. Backup Exec's data lifecycle management feature automatically deletes expired backup sets and reclaims the disk space. If you want to keep the backup data longer than the period that you specify when you create the backup job, you should create a duplicate backup job. A duplicate backup job can copy the backup data from the original storage device to tape media or to disk cartridge media, which you can then send for

long-term or off-site storage. You can also keep the backup sets from automatically expiring by retaining the backup sets. Backup Exec then retains all dependent backup sets as well.

To be eligible for configuration as disk storage, a disk must have at least 1 GB of disk space and cannot be configured as deduplication disk storage. Although you can configure disk storage and deduplication disk storage on the same disk, it is not recommended.

When you create disk storage on a disk that is attached to the network, you must specify the path to an existing share. You should use the server name in the UNC path rather than an IP address.

Note: Before you create the disk storage on a network share, you must give read and write permissions to the Backup Exec service account. The Backup Exec service account is on the Backup Exec server that you want to access the network share.

For information about the best practices to use hot-pluggable devices in Backup Exec, refer to *Backup Exec Best Practices*.

When you create disk storage on a local disk, Backup Exec lets you specify any of the following locations:

- Volumes with or without drive letters.
You can create only one disk storage on a volume.
- Unformatted partitions.
Backup Exec formats and partitions the drive for you, if necessary.
- Drives that do not have partitions.

Backup Exec creates a folder named BEControl on the root of the volume. Do not delete or edit the contents of the BEControl folder, and do not copy it to other volumes or drive letters.

In Windows Explorer, the backup files that the disk storage device contains display with a .bkf file extension. Each disk storage device also contains a file named changer.cfg and a file named folder.cfg, which store information about the backup files. Do not delete or edit the changer.cfg or folder.cfg files.

A subfolder with a prefix of IMG in the name may display in a disk storage device. This subfolder appears if the option to enable Granular Recovery Technology (GRT) was selected for backup, or if you select the disk storage device as storage for backup data.

You must use the **Configure Storage** wizard to create disk storage. In the **Configure Storage** wizard, Backup Exec provides a list of disks on which you can

create disk storage. The disks do not appear in the list in the alphabetical order of the drive letter. Instead, the disk that appears first in the list has the most amount of disk space. You can select any disk that you want, but the disk that Backup Exec recommends for use appears at the top of the list. The disk that you use as the system drive always appears last in the list. It is recommended that you do not configure disk storage on the system drive.

Note: If Windows data deduplication is enabled on the disk storage volume, Backup Exec excludes the backup data in the folder \BEData from deduplication, unless the \BEData folder already exists. Backup Exec must exclude backup data from deduplication for you to use Simplified Disaster Recovery (SDR) to perform a local recovery of the Backup Exec server.

If Windows data deduplication is enabled on the disk storage volume, local disaster recovery using SDR fails. The Windows Preinstallation Environment (Windows PE) that SDR uses cannot read the files that Windows data deduplication processes.

To configure disk storage

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server feature is not installed	Click Disk-based storage , and then click Next .
If the Central Admin Server feature is installed	Do the following in the order listed: ■ Select the Backup Exec server on which you want to configure storage, and then click Next. ■ Click Disk-based storage, and then click Next.
- 3 Click **Disk storage**, and then click **Next**.
- 4 Enter a name and description for the disk storage device, and then click **Next**.
- 5 Specify if you want to create the disk storage device on a local disk or on a network share, enter the location or path, and then click **Next**.
- 6 Specify how many write operations to let run at the same time on this disk storage device, and then click **Next**.
- 7 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the disk storage device

Click **Finish**.

See [“Changing the location of a disk storage device”](#) on page 337.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Changing the location of a disk storage device

You can change the location of an existing disk storage device. You must have a different volume available to which you can move the files in the \BEData folder.

Note: When you copy files from the original disk storage device to the new location, do not copy .cfg files.

To change the location of a disk storage device

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Click **Disk storage**, and then click **Next**.
- 3 Enter a different name and description than the original disk storage device, and then click **Next**.
- 4 Specify a different drive letter than the original disk storage device, and then click **Next**.
- 5 Specify the number of write operations that you want to let run at the same time on this disk storage device, and then click **Next**.
- 6 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the disk storage device Click **Finish**.

- 7 In Windows Explorer, copy and paste the following files from the \BEData folder on the original volume to the \BEData folder on the new volume:
 - .Bkf files
 - Any subfolders with a prefix of IMG in the name
- 8 In Windows Explorer, delete all of the files from the original disk storage device.
- 9 On the Backup Exec Administration Console, on the **Storage** tab, right-click the original disk storage device, and then click **Delete**.
- 10 Rename the new disk storage device with the name of the original disk storage device.
- 11 Right-click the new disk storage device, and then click **Inventory and Catalog**.

See [“Inventorying and cataloging a storage device”](#) on page 563.

See [“Configuring disk storage”](#) on page 334.

Editing disk storage properties

You can edit disk space management settings for the disk storage device.

To edit disk storage properties

- 1 On the **Storage** tab, double-click the storage for which you want to edit properties.
- 2 In the left pane, click **Properties**.
- 3 Edit any of the following options:

Name	Displays the name of the disk storage. You can edit this field.
Description	Displays a description of the disk storage. You can edit this field.

Limit Backup Exec to read-only operations

Prevents Backup Exec from deleting expired backup sets on this disk storage when you reattach the disk storage to the Backup Exec server. Otherwise, Backup Exec's data lifecycle management feature deletes any backup sets that are expired and reclaims the disk space.

The default value is **No**.

This option applies only when the disk storage has been detached from the Backup Exec server for the number of days that you specify in the global setting. The default number of days is 14.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

See [“Backup sets ”](#) on page 359.

See [“How to restore data from a reattached or reinserted disk-based storage device”](#) on page 344.

Maximum file size

Displays the maximum file size on the disk storage. The data from the backup job is contained in a file on the disk.

The default value is 50 GB or the capacity of the disk storage.

Preallocate disk space incrementally up to the maximum file size	Creates the file when the backup job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job used. For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB. The default value is Disabled.
Preallocation increment	Displays the amount of disk space by which to increase the file size. The file size increases by this increment as the job requires disk space, up to the maximum file size. The default value is 1 GB.
Auto detect block and buffer size	Indicates if Backup Exec automatically detects the preferred settings for the block size and buffer size for the disk storage. The default value is Enabled. If you disable this setting, you can then choose the block size and buffer size to use.

Block size

Displays the size of the blocks of data that are written to new media in this disk storage device if the option **Auto detect block and buffer size** is disabled. The default is the preferred block size.

Some storage devices provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a storage device that supports larger block sizes, you can change the block size. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.

See the manufacturer's documentation for help in configuring the device.

Backup Exec does not ensure that the storage device supports the requested block size. If the requested block size is not supported, it defaults to its standard block size.

If the device does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of the data that is sent to the disk storage device on each read or write request if the option **Auto detect block and buffer size** is disabled. The buffer size must be an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve storage performance. Each type of storage device requires a different buffer size to achieve maximum throughput.

If the preferred block size is greater than 64 KB, the default buffer size is the same as the default block size. If the preferred block size is less than 64 KB, then the default buffer size is 64 KB.

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns red to indicate critically low available space. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold. This threshold must be less than the warning threshold.

The default value is 5%.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns orange to indicate a low disk space condition. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold. This threshold must be less than the low disk space threshold.

The default value is 15%.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns yellow to indicate the first of three low disk space conditions. If free disk space drops below this threshold to the amount that is specified in the warning threshold, another alert is sent. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of disk space does not include the disk space that is reserved for non-Backup Exec operations.

When low disk space reaches this threshold, data lifecycle management immediately searches the device for expired backup sets that it can delete.

You can change the value of the threshold.

The default value is 25%.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Disk space to reserve for non-Backup Exec operations

Displays the amount of disk space to set aside for applications other than Backup Exec.

The default value is 10 MB.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the disk storage.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk storage device.
- You want this disk storage to allow buffered read, which is the reading of large blocks of data.

Enabling buffered reads may provide increased performance.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk storage device.
- You want this disk storage to allow buffered write, which is the writing of large blocks of data.

Concurrent write sessions

Displays the number of concurrent write operations that you want to allow to this disk storage device.

4 Click **Apply**.

See [“Configuring disk storage”](#) on page 334.

How to restore data from a reattached or reinserted disk-based storage device

If the backup sets on a disk storage device or a disk cartridge expire while that device is detached, Backup Exec deletes the catalogs for those backup sets. To restore from those backup sets at some future time, you must run an inventory and catalog operation on the device when you reattach it. When you run the inventory and catalog operation, Backup Exec sets a new expiration date for each backup set by using the backup set's original retention setting, calculated from the current date. Backup Exec also resets the expiration date for any backup set on the storage device that expires within seven days of the current date.

If you want the backup sets to expire, you can disable the storage device property **Limit Backup Exec to read-only operations**. To find this option, on the **Storage** tab, right-click the storage device, click **Details**, and then click **Properties**. Do not run an inventory and catalog operation. Backup Exec reclaims the disk space on that storage device during data lifecycle management. You can also delete the backup sets.

See [“Inventorying and cataloging a storage device”](#) on page 563.

See [“Backup sets ”](#) on page 359.

Configuring disk cartridge storage

Disk cartridges are a type of storage that usually remains attached to the Backup Exec server while you remove the media, such as RDX. If you are not sure if the storage has removable media, you can open the Computer folder on your Windows computer. The devices that contain removable media are listed.

Backup Exec uses data lifecycle management to automatically expire the backup sets that are stored on disk cartridge media. The backup sets on disk cartridge media are kept for the amount of time that you specify in the backup job properties. Backup Exec automatically reclaims the space as the backup data expires. You can keep the backup sets from automatically expiring by retaining the backup sets. Backup Exec then retains all dependent backup sets as well.

You must use the **Configure Storage** wizard to configure a disk cartridge device.

Available storage operations for disk cartridges and media are dependent on the type of disk cartridge that you have. For example, fewer operations are available for USB memory sticks than are available for RDX devices.

To configure disk cartridge storage

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server feature is not installed	Click Disk-based storage , and then click Next .

If the Central Admin Server feature is installed	Do the following in the order listed: ■ Select the Backup Exec server that you want to configure storage for, and then click Next. ■ Click Disk-based storage, and then click Next.

- 3 Click **Disk cartridge device**, and then click **Next**.
- 4 Enter a name and description for the disk cartridge device, and then click **Next**.
- 5 Specify where the disk cartridge device is located, and then click **Next**.
- 6 Review the summary, and then do one of the following:

- To change the configuration
- Do the following in the order listed:
- Click the heading that contains the items that you want to change.

■ Make any changes, and then click **Next** until the summary appears.

■ Click **Finish**.
- To configure the disk cartridge device
- Click **Finish**.

See [“Editing disk cartridge properties”](#) on page 346.

Editing disk cartridge properties

You can edit settings for the disk cartridge storage.

See [“Configuring disk cartridge storage”](#) on page 345.

To edit disk cartridge properties

- 1 On the **Storage** tab, double-click the disk cartridge for which you want to edit properties.
- 2 In the left pane, click **Properties**.
- 3 Edit any of the following options:

Name	Displays the name of the disk cartridge. Disk cartridge names cannot exceed 128 characters. You can rename the disk cartridge.
Description	Displays a description of the disk cartridge. You can change the description.
Maximum file size	Displays the maximum file size on the disk cartridge. The data from the job is contained in a file on the disk cartridge. The default value is 50 GB or the capacity of the disk cartridge media.

Preallocate disk space incrementally up to the maximum file size	Creates the file when the job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job used. For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB. The default value is Disabled.
Preallocation increment	Displays the amount of disk space by which to increase the file size if the option to preallocate disk space is enabled. The file size increases by this increment as the job requires disk space, up to the maximum file size. The default value is 1 GB.
Auto detect block and buffer size	Indicates if Backup Exec automatically detects the preferred settings for the block size and buffer size for the disk storage. The default value is Enabled. If you disable this setting, you can then choose the block size and buffer size to use.

Block size

Displays the size of the blocks of data that are written to new media in this disk cartridge if the option **Auto detect block and buffer size** is disabled. The default is the preferred block size.

Some storage devices provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use the storage that supports larger block sizes, you can change the block size. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.

See the manufacturer's documentation for help in configuring the storage.

Backup Exec does not ensure that the storage device supports the requested block size. If the requested block size is not supported, it defaults to its standard block size.

If the storage does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of the data that is sent to the disk cartridge on each read or write request if the option **Auto detect block and buffer size** is disabled. The buffer size must be an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve storage performance. Each type of storage requires a different buffer size to achieve maximum throughput.

If the preferred block size is greater than 64 KB, the default buffer size is the same as the default block size. If the preferred block size is less than 64 KB, then the default buffer size is 64 KB.

Low disk space - Critical

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns red to indicate critically low available space. Backup Exec does not send low disk space alerts for disk cartridge devices.

You can change the value of the threshold, but it must be less than the warning threshold.

The default value is 5%.

See [“Storage tab overview in Backup Exec”](#) on page 535.

Low disk space - Warning

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns orange to indicate a low disk space condition. Backup Exec does not send low disk space alerts for disk cartridge devices.

You can change the value of the threshold, but it must be less than the low disk space threshold.

The default value is 15%.

See [“Storage tab overview in Backup Exec”](#) on page 535.

Low disk space

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns yellow to indicate the first of three low disk space conditions. Backup Exec does not send low disk space alerts for disk cartridge devices. When the disk cartridge media reaches this threshold, the data lifecycle management feature immediately searches this disk cartridge media for expired backup sets that it can delete

You can change the value of the threshold.

The default value is 25%.

See [“Storage tab overview in Backup Exec”](#) on page 535.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the disk cartridge.

The default value is **Enabled**.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk cartridge.
- You want this disk cartridge to allow buffered read, which is the reading of large blocks of data.

Enabling buffered read operations may provide increased performance.

The default value is **Enabled**. If you disable **Auto detect settings**, this setting also changes to **Disabled**.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk cartridge.
- You want this disk cartridge to allow buffered write, which is the writing of large blocks of data.

The default value is **Enabled**. If you disable **Auto detect settings**, this setting also changes to **Disabled**.

4 Click **Apply.**

Editing disk cartridge media properties

You can edit the properties of the disk cartridge media.

See [“Configuring disk cartridge storage”](#) on page 345.

To edit disk cartridge media properties

- 1 On the **Storage** tab, double-click the disk cartridge device that contains the media.
- 2 In the left pane, click **Media**.

3 Edit any of the following options:

Media label

Displays the media label that Backup Exec assigns automatically or that the administrator assigns.

You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the Backup Exec user interface.

Media description

Displays the original media label if the media is imported media. You can edit the media description to make it a more descriptive label. The description is limited to 128 characters.

Preserve description

Keeps the media description when you select **Yes** in the drop-down box. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

Limit Backup Exec to read-only operations

Prevents Backup Exec from deleting expired backup sets on this disk cartridge media when you reinsert the cartridge media into the storage device. If you select **Yes** in the drop-down menu, then Backup Exec's data lifecycle management feature deletes any backup sets that are expired and reclaims the disk space.

The default value is **No**.

This option applies only when the disk cartridge media is absent from the Backup Exec server for the number of days that you specify in the global setting. By default, the number of days that the disk cartridge media must be absent before this option takes effect is 30 days.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

See [“Backup sets ”](#) on page 359.

See [“How to restore data from a reattached or reinserted disk-based storage device”](#) on page 344.

4 Click **Apply**.

How data lifecycle management (DLM) deletes expired backup sets on disk-based storage

Backup Exec uses data lifecycle management (DLM) to automatically delete expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, virtual disks, and cloud storage. You specify how long to keep backup data when you create a backup job that is sent to a disk-based storage device. When the amount of time to keep the backup data expires, the data lifecycle management feature deletes the backup sets and reclaims the disk space unless there are dependent backup sets such as incrementals.

By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. If backup sets are dependent on other backup sets, then Backup Exec does not delete the backup set until all expiration dates on the backup sets are reached.

Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

For example, you create a backup definition that contains a full backup and an incremental backup for the C: volume on a server. The first full backup runs, followed by the first incremental backup, and then the second incremental backup. The second full backup runs, followed by the third incremental backup, and then the fourth incremental backup. No more backups are run. All of the backup sets created by these backups eventually expire and are deleted by DLM. However, the backup sets that resulted from the second full backup and the third and fourth incremental backups are kept.

Backup Exec keeps these related backup sets because they are the most recent backup sets that you need to restore the C: volume. By keeping the last related backup sets, you have the data to restore the volume.

Warning: DLM deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup.

To prevent the backup sets from being automatically deleted, you can manually retain specific backup sets or you can change the expiration date of the backup set. If you retain a backup set, Backup Exec then retains all dependent backup sets as well.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 363.

See [“Changing the expiration date of backup sets on disk-based storage”](#) on page 362.

DLM searches for expired backup sets to delete from disk-based storage at the following times:

- Every hour.
DLM starts for the first time 1 hour after you install Backup Exec and the Backup Exec service starts, and then runs every hour after that. If you restart the Backup Exec service, the hourly DLM cycle also restarts.
- When the low disk space threshold for disk-based storage is reached.
The low disk space threshold is a storage device property. When the storage device's used capacity reaches this threshold, DLM immediately searches the device for the expired backup sets that it can delete.
- When you manually expire a backup set.
When you manually expire a backup set on a standalone Backup Exec server, DLM runs immediately on the storage device on which the backup set is

located. In a Central Admin Server feature (CAS) environment, if you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired.

To monitor the backup sets that data lifecycle management deletes, you can view the **Backup Set Retention** category in the audit log. You can also run the audit log report to view the backup sets that data lifecycle management deletes.

See [“Configuring audit logs”](#) on page 774.

See [“Audit Log report”](#) on page 810.

Storage options that can affect how data lifecycle management deletes backup sets are described in the following table:

Table 8-4 Storage options in Backup Exec that affect data lifecycle management

Storage option	Description
Allow Backup Exec to delete all expired backup sets	

Table 8-4 Storage options in Backup Exec that affect data lifecycle management *(continued)*

Storage option	Description
	This global setting lets Backup Exec delete expired backup sets, even if they are the last remaining backup sets that you need to restore a server. Use this option when you want to delete backup data after a period of time. Alternatively, you can manually delete backup sets. Warning: If you enable this option, the data that you need to restore a server may not be available. The following conditions may cause you to lose backup sets when you enable this option: ■ If the backup sets from the last full backup job expire before the next full backup runs. Ensure that when you create jobs, the backup sets are kept longer than the amount of time between full backups. ■ If the backup job fails or is missed, and is not rerun before the backup sets expire. Monitor any failed or missed jobs, and ensure that you rerun them before the backup sets from the previous full backup expire. Note: In a Central Admin Server feature (CAS) environment, this option is only available on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on all of the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers in a CAS environment. To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage.

Table 8-4 Storage options in Backup Exec that affect data lifecycle management *(continued)*

Storage option	Description
	See “Editing global settings for storage” on page 548.
Limit Backup Exec to read-only operations on a disk-based storage device if it has been detached for x number of days	This global setting prevents Backup Exec from deleting expired backup sets on any disk-based storage device as soon as you reattach it. When you limit Backup Exec to read-only operations, you have time to review any expired backup sets and determine if you want to keep them before data lifecycle management deletes them. To keep expired backup sets, you can retain them or change the expiration date. See “Retaining backup sets on disk-based storage to prevent them from expiring” on page 363. See “Changing the expiration date of backup sets on disk-based storage” on page 362. Backup Exec is limited to read-only operations only on the disk-based storage devices that are detached from the Backup Exec server for the specified number of days. To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage. See “Editing global settings for storage” on page 548.

Table 8-4 Storage options in Backup Exec that affect data lifecycle management *(continued)*

Storage option	Description
Limit Backup Exec to read-only operations on a disk cartridge if it has not been inserted for x number of days	This global setting prevents Backup Exec from deleting expired backup sets on any disk cartridge as soon as you insert it into a disk cartridge device. When you limit Backup Exec to read-only operations, you have time to review any expired backup sets and determine if you want to keep them before data lifecycle management deletes them. To keep expired backup sets, you can retain them or change the expiration date. See “Backup sets” on page 359. Backup Exec is limited to read-only operations only on the disk cartridges that are detached from the Backup Exec server for the specified number of days. To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage. See “Editing global settings for storage” on page 548.
Limit Backup Exec to read-only operations	This device property prevents data lifecycle management from running on a specific disk-based or disk cartridge storage device that you reattach or reinsert it. This option applies only when the storage device has been absent from the Backup Exec server for the number of days that you specify in either of the two previous global settings described in this table. To access this option, on the Storage tab, right-click the device, click Details, and then click Properties. See “Editing disk storage properties” on page 338. See “Editing disk cartridge properties” on page 346. See “Editing the properties of a deduplication disk storage device” on page 985.

Table 8-4 Storage options in Backup Exec that affect data lifecycle management (*continued*)

Storage option	Description
Low disk space	This device property is the first of three low disk space conditions. When the storage device's used capacity reaches this threshold, the data lifecycle management feature immediately searches the device for expired backup sets that it can delete. Data lifecycle management also runs on a disk cartridge if the cartridge reaches capacity during a backup job. The backup may not need to span to another cartridge if data lifecycle management deletes enough expired backup sets. If a job does span and you insert a new disk cartridge that is full, data lifecycle management deletes expired backup sets on the new cartridge. To access this option, on the Storage tab, right-click the device, click Details, and then click Properties. See "Setting low disk space thresholds on disk-based storage" on page 333.

For information about the best practices of Backup Exec data lifecycle management (DLM) feature, refer to *Backup Exec Best Practices*.

Backup sets

A backup set is a collection of the data that you back up from a single source of content. A single source of content can be a server or a Microsoft Exchange data set, for example. If you select multiple sources of content, Backup Exec creates multiple backup sets. When you run a backup job, Backup Exec creates the backup sets and writes them on storage. To restore data, you select the backup sets that contain the data that you want to restore.

Backup Exec keeps the backup sets that are stored on disk storage and disk cartridge media for as long as you specify in the backup job properties. By default, the amount of time that backup sets are stored is based on the type of backup job and its schedule.

For example, you can specify to keep the backup sets from a full backup for two weeks on a disk-based storage device. After two weeks, the backup sets expire

and Backup Exec uses the data lifecycle management feature to delete the backup sets and reclaim that disk space. If you later create an incremental backup job, Backup Exec keeps the full backup sets for two weeks, plus the amount of time that it keeps the incremental backup sets. If you keep the incremental backup sets for four weeks, then Backup Exec keeps the full backup sets for six weeks. The data from a full backup job is kept as long as the data from its associated incremental backup jobs. Backup Exec does not reclaim the disk space for backup sets from a job that depends on another job until the data retention expires for all of the associated jobs. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

Backup Exec manages the retention of backup sets differently depending on the type of storage to which you back up the data.

Table 8-5 Storage types and backup set retention

Type of storage	Backup data retention
Disk storage, disk cartridge devices, deduplication disk storage, storage arrays, cloud storage, and virtual disks	Backup Exec uses data lifecycle management to automatically delete expired backup sets from disk-based storage. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well. For disk-based storage and any disk cartridges that you reattach to the Backup Exec server after a specified number of days, you can prevent Backup Exec from reclaiming that disk space. A global setting limits Backup Exec to read-only operations on disk-based storage or a disk cartridge if it has been detached for a specified number of days. You can also limit Backup Exec to read-only operations per disk storage or disk cartridge by enabling the setting on the device properties. See “How data lifecycle management (DLM) deletes expired backup sets on disk-based storage” on page 352. See “How to restore data from a reattached or reinserted disk-based storage device” on page 344.

Table 8-5 Storage types and backup set retention (*continued*)

Type of storage	Backup data retention
Tape cartridge media	Backup Exec uses the Advanced Device and Media Management (ADAMM) feature to manage data retention on tape cartridge media. ADAMM expires the backup sets that are stored on media according to a set of rules that you apply to the media. Backup sets are not automatically deleted from tape cartridge media, but they can be overwritten, depending on the rules that you specify. The set of rules that manages tape cartridge media is called a media set. You create media sets that specify append periods, overwrite protection periods, and vaulting periods. See “Default media sets” on page 491.

You can perform the following actions on backup sets that are on disk-based storage:

- Change the expiration date of backup sets to keep them longer or expire them immediately.
- Extend the amount of time that you keep backup sets by retaining them.
- Release any backup sets that are retained to let them expire automatically.

For all backup sets, including those on tape cartridge media, you can perform the following actions:

- Catalog backup sets so that you can view the data that is contained in the backup sets and search for files to restore.
- View the contents of backup sets and browse the backed up data that is contained in them.
- View the system properties and job properties of backup sets.

See [“Changing the expiration date of backup sets on disk-based storage”](#) on page 362.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 363.

See [“Releasing retained backup sets on disk-based storage”](#) on page 364.

See [“Cataloging backup sets”](#) on page 260.

See [“Viewing the contents or properties of backup sets”](#) on page 365.

Changing the expiration date of backup sets on disk-based storage

You can change the expiration date of backup sets that are on disk-based storage to keep them for a longer or shorter amount of time. You can also expire backup sets immediately, if you no longer want to keep them. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

When you manually expire backup sets, Backup Exec checks those backup sets to ensure that no other backup sets are dependent upon them. Backup sets from incremental and differential jobs are dependent upon the backup sets that come from the full backup job in the same backup definition. You cannot expire only the backup sets that come from the full backup job because the dependent backup sets would not function without them. When Backup Exec detects dependent backup sets, it gives you the option to expire the backup sets and any dependent backup sets.

When you manually expire a backup set on a standalone Backup Exec server, data lifecycle management (DLM) runs immediately on the storage device on which the backup set is located and deletes the expired backup set. In a Central Admin Server feature (CAS) environment, if you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created and deletes the expired backup set. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired to delete the expired backup set. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired and deletes the expired backup set.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

To change the expiration date of backup sets on disk-based storage

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To change the expiration date for a single backup set, right-click the backup set.
 - To change the expiration date for multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Do either of the following:

To expire backup sets immediately

- Click **Expire**.
Backup Exec displays the backup set and any dependent backup sets.
- To expire the single backup you selected and any dependent backup sets, click **Expire**.
If you selected multiple backup sets, click **Expire** or **Expire All**. You can also click **Skip** to skip a backup set and its dependents from deletion.

To change the expiration date of backup sets

- Click **Expiration**.
- Enter the new expiration date in the **Expiration** field.
- Click **OK**.

See [“Backup sets”](#) on page 359.

Retaining backup sets on disk-based storage to prevent them from expiring

You can prevent backup sets on disk-based storage from automatically expiring by retaining the backup sets. Backup Exec retains all dependent backup sets as well. For example, if you choose to retain an incremental backup set, Backup Exec retains all backup sets dating back to, and including, the last full backup job. You may need to retain backup sets for legal purposes, such as compliance with data retention laws.

After you retain a backup set, Backup Exec prevents the backup set from expiring indefinitely. If you decide that you no longer need to retain a backup set, you must release it so that it can expire automatically. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

To retain backup sets on disk-based storage to prevent them from expiring

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to retain.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To retain a single backup set, right-click the backup set.

- To retain multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Retain**.
 - 5 In the **Reason to retain backup sets** field, select the reason that you want to retain the backup sets. You can choose from the following options:

Legal	Select this option if the reason for retaining the backup sets is a legal one. You may have to retain backup sets to comply with corporate or regulatory data retention policies.
User defined	Select this option if the reason for retaining the backup sets is something other than a legal one.
System defined	This option is used by Backup Exec when a backup set needs to be retained for some future operations to run. This option is grayed out and cannot be selected.
 - 6 In the **Explanation** field, type any additional information about why you retained the backup sets. Entering an explanation in this field can help remind you why you retained the backup sets or for how long they should be retained.
 - 7 Click **OK**.

See [“Backup sets”](#) on page 359.

See [“Releasing retained backup sets on disk-based storage”](#) on page 364.

Releasing retained backup sets on disk-based storage

You can override the retention period for backup sets that are on disk-based storage by manually retaining them. When you choose to retain backup sets on disk-based storage, Backup Exec prevents the backup sets from automatically expiring when their retention period is over. You can manually retain backup sets indefinitely.

If you no longer need the retained backup sets, you can allow them to expire. First, you need to remove the backup sets' retained status. Then Backup Exec expires the backup sets automatically according to the backup sets' storage settings. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

To release retained backup sets on disk-based storage

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to release.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To release a single backup set, right-click the backup set.
 - To release multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Retain**.
- 5 Select **Do not retain**.
- 6 Click **OK**.

See [“Backup sets ”](#) on page 359.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 363.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Viewing the contents or properties of backup sets

After you complete a backup job, you can view the data that is contained in the backup sets that are created. Viewing the contents of backup sets can help you to confirm what data was backed up. You may also want to view the contents of backup sets before you run a restore job to verify the data that they contain.

You can also view the following backup set properties:

- Backup source
- Backup date
- Expiration date
- Backup method
- Size
- Location
- Backup set description
- Data encryption
- True image

- Server name
- Catalog file name
- Snapshot

To view the contents or properties of backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to view.
- 2 In the left pane, click **Backup Sets**.
- 3 Double-click the backup set that you want to view.

Note: On the **Backup and Restore** tab, you must expand the backup source to see the backup set.

- 4 Do either of the following:

To view the contents of the backup sets In the left pane, click **Contents**.

The contents of the backup set display in the left pane in a tree view. You can expand folders and drives to view their contents in the right pane.

To view the properties of the backup sets In the left pane, click **Properties**.

See [“Backup sets ”](#) on page 359.

Cloud-based storage devices

This chapter includes the following topics:

- [About cloud-based storage devices](#)
- [Amazon S3 cloud-based storage](#)
- [Google cloud-based storage](#)
- [Microsoft Azure cloud-based storage](#)
- [Private cloud-based storage](#)
- [About S3-Compatible Cloud Storage](#)
- [About Backup Exec Cloud Deduplication](#)
- [Notes for Backup Exec Cloud Deduplication](#)
- [Cloud deduplication storage device](#)
- [Notes for cloud-based storage devices](#)
- [Editing the properties of a cloud-based storage device](#)
- [Best practices for using cloud-based storage](#)
- [Changing default cloud storage settings](#)
- [About the Backup Exec™ CloudConnect Optimizer](#)

About cloud-based storage devices

Backup Exec supports backups to cloud-based storage devices. You can use the cloud connector to back up and restore data from cloud Storage as a Service (STaaS) vendors. Cloud-based storage is unlike traditional tape or disk media, which use persistent backup images. Usually, a public cloud storage vendor calculates cloud-based storage costs per byte stored and per byte transferred.

For the list of supported public and private cloud providers, refer to the Backup Exec hardware compatibility list.

Note: A cloud storage device cannot belong to any storage pools.

For information about the best practices to manage Backup Exec Cloud Connector, refer to *Backup Exec Best Practices*.

The following notes apply to cloud-based storage devices:

- If you use Backup Exec Central Admin Server feature, you can share a public cloud storage device between multiple managed Backup Exec servers. You can enable the sharing when you add a public cloud storage device. You can select new managed Backup Exec servers to share a public cloud storage device. You can remove the sharing ability for any managed Backup Exec servers at any time.
- Data lifecycle management automatically expires the backup sets that are on cloud storage.
- Some cloud storage providers require encryption.
See [“Using encryption with Backup Exec”](#) on page 729.

Amazon S3 cloud-based storage

The following sections provide information about the requirements for configuring an Amazon S3 cloud-based storage device and configure the storage for Amazon S3 storage in Backup Exec.

See [“Requirements for configuring an Amazon S3 cloud-based storage device”](#) on page 369.

See [“Configuring storage for Amazon cloud storage”](#) on page 369.

Requirements for configuring an Amazon S3 cloud-based storage device

Backup Exec cloud connector enables Backup Exec to back up data to and restore data from Amazon Simple Storage Service (S3).

Review the following requirements before configuring an Amazon S3 cloud-based storage device:

- You must configure buckets using S3 service from your AWS account. You must copy the Amazon access key and secret access key from the Amazon portal.
- Ensure that you have already created the S3 buckets. Buckets represent a logical unit of storage on the cloud-based storage device.

Note: As a best practice, you should create specific buckets to use exclusively with Backup Exec.

Each cloud storage device must use a different bucket. Do not use the same bucket for multiple cloud storage devices even if these devices are configured on different Backup Exec servers.

- Ensure that the bucket names meet the following Backup Exec requirements:
 - Bucket names can contain lowercase letters, numbers, and dashes (or hyphens)
 - Bucket names cannot begin with a dash (or a hyphen)

The buckets are not available for use in Backup Exec if the bucket name does not comply with the bucket naming convention or if you have created a bucket for a region that Backup Exec does not support.

To review the list of supported regions, refer to the Backup Exec hardware compatibility list.

See [“Configuring storage for Amazon cloud storage”](#) on page 369.

Configuring storage for Amazon cloud storage

You can configure a cloud-based storage device for the Amazon cloud storage, and then back up data to it.

Note: For cloud-based storage devices, by default, the **Do not verify data for this job** option is now selected in the **Backup Options**. Cloud vendors charge for operations that read data from and write data to the cloud. To avoid charges for reading data during the verify operation of a backup or duplicate job, this option is selected by default.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Requirements for configuring an Amazon S3 cloud-based storage device”](#) on page 369.

To configure storage for Amazon cloud storage in Backup Exec

- 1 On the **Backup and Restore** tab, click the **Back Up > Backup to Cloud** option and then click **OK**.

Or, on the **Storage** tab, in the **Configure Storage** group, click **Cloud Storage** and then click **Next**.

Or, on the **Storage** tab, click **Configure Cloud Storage**.
- 2 Enter a name and description for the cloud storage device, and then click **Next**.

Note: If you want to use the WORM enabled cloud deduplication storage to create backup sets, ensure that you enter the cloud volume name that you created when you ran `msdpclutil.exe`.

- 3 From the list of cloud storage providers, select **S3**.
- 4 (Optional) Click the link to add a new S3 cloud instance using the Generic S3 Configurator tool.
- 5 Click **Next**.
- 6 In the **Cloud storage** field, select the Amazon server name from the drop-down list.

- 7 In the **Storage tier** field, select a tier from the drop-down list.

The options are **Standard** (for frequently accessed data), **Standard_IA** (Infrequently accessed data), **Glacier**, **Deep Archive**, **Intelligent_Tiering**, or **One Zone_IA** (Infrequently accessed data).

Note: To restore from **Glacier** and **Deep Archive** storage device is a time consuming activity. Refer to the Amazon documentation for more information about retrieval time delay.

For more information about the storage tiers, refer to the following link:

<https://aws.amazon.com/s3/storage-classes/>

- 8 In the **Logon account** field, select an account from the drop-down list or click **Add/Edit** to add an account.
- 9 On the **Logon account selection** dialog box, click **Add**.
- 10 On the **Add Logon Credentials** dialog box, do the following:
 - In the **User name** field, type the Amazon account access key ID.
 - In the **Password** field, type the Amazon account secret access key.
 - In the **Confirm password** field, type the Amazon account secret access key again.
 - In the **Account name** field, type a name for this logon account.
The Backup Exec user interface displays this name as the cloud storage device name in all storage device options lists.
- 11 Click **OK** twice.
- 12 Select the Amazon logon account that you created in step 7, and then click **Next**.
- 13 Select a bucket from the list of buckets that are associated with the server name and then click **Next**.

Note: If you want to use the WORM enabled cloud deduplication storage to create backup sets, ensure that you select the bucket name that you created when you ran `msdpclutil.exe`.

- 14 Specify how many concurrent operations can run at the same time on this cloud storage device, and then click **Next**.

This setting determines the number of jobs that can run at the same time on this device. The suitable value for this setting may vary depending on your environment and the bandwidth to the cloud storage. You may choose the default value.

- 15 Review the configuration summary, and then click **Finish**.

Backup Exec creates a cloud storage device. You must restart Backup Exec services to bring the new device online.

- 16 In the window that prompts you to restart the Backup Exec services, click **Yes**.

After services restart, Backup Exec displays the new cloud storage location in the **All Storage** list.

Google cloud-based storage

The following sections provide information about the requirements to configure a Google cloud-based storage device and configure the storage for Google storage in Backup Exec.

See [“Requirements for configuring a Google cloud-based storage device”](#) on page 372.

See [“Configuring storage for Google cloud storage”](#) on page 373.

Requirements for configuring a Google cloud-based storage device

Backup Exec cloud connector enables Backup Exec to back up data to and restore data from Google cloud storage.

Review the following requirements before configuring a Google cloud-based storage device:

- You must configure a Google cloud platform account. You must copy the Google S3 interoperability access key and a Google S3 interoperability secret key from the Google portal.
You can generate these keys in the Google Developers Console.
- Ensure that you have already created the buckets. Buckets represent a logical unit of storage on the cloud-based storage device.
Backup Exec supports all location types from Google - Region, Dual-region and Multi-region.

Note: As a best practice, you should create specific buckets to use exclusively with Backup Exec.

Each cloud storage device must use a different bucket. Do not use the same bucket for multiple cloud storage devices even if these devices are configured on different Backup Exec servers.

- Ensure that the bucket names meet the following Backup Exec requirements:

- Bucket names can contain lowercase letters, numbers, and dashes (or hyphens)

- Bucket names cannot begin with a dash (or a hyphen)

The buckets are not available for use in Backup Exec if the bucket name does not comply with the bucket naming convention or if you have created a bucket for a region that Backup Exec does not support.

To review the list of supported regions, refer to the Backup Exec hardware compatibility list.

See [“Configuring storage for Google cloud storage”](#) on page 373.

Configuring storage for Google cloud storage

You can configure a cloud-based storage device for the Google cloud storage, and then back up data to it.

Note: For cloud-based storage devices, by default, the **Do not verify data for this job** option is now selected in the **Backup Options**. Cloud vendors charge for operations that read data from and write data to the cloud. To avoid charges for reading data during the verify operation of a backup or duplicate job, this option is selected by default.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Requirements for configuring a Google cloud-based storage device”](#) on page 372.

To configure storage for Google cloud storage in Backup Exec

- 1 On the **Backup and Restore** tab, click the **Back Up > Backup to Cloud** option and then click **OK**.

Or, on the **Storage** tab, in the **Configure Storage** group, click **Cloud Storage** and then click **Next**.

Or, on the **Storage** tab, click **Configure Cloud Storage**.

- 2 Enter a name and description for the cloud storage device, and then click **Next**.
- 3 From the list of cloud storage providers, select **S3**, and then click **Next**.
- 4 In the **Cloud storage** field, select the Google server name from the drop-down list.
- 5 In the **Storage tier** field, select a tier from the drop-down list. The options are **Standard** (for frequently accessed data), **Nearline** (infrequently accessed data), **Coldline** (infrequently accessed data), **Archive**.

For more information about the storage tiers, refer to the following link:

<https://cloud.google.com/storage/docs/storage-classes>

- 6 In the **Logon account** field, select an account from the drop-down list or click **Add/Edit** to add an account.
- 7 On the **Logon Account Selection** dialog box, click **Add**.
- 8 On the **Add Logon Credentials** dialog box, do the following:
 - In the **User name** field, type the Google account access key ID.
 - In the **Password** field, type the Google account secret access key.
 - In the **Confirm password** field, type the Google account secret access key again.
 - In the **Account name** field, type a name for this logon account.

The Backup Exec user interface displays this name as the cloud storage device name in all storage device options lists.

- 9 Click **OK** twice.
- 10 Select the Google logon account that you created in step 7, and then click **Next**.

Backup Exec supports all location types from Google - Region, Dual-region, and Multi-region.

- 11 Select a bucket from the list of buckets that are associated with the server name and then click **Next**.

- 12 Specify how many concurrent operations can run at the same time on this cloud storage device, and then click **Next**.

This setting determines the number of jobs that can run at the same time on this device. The suitable value for this setting may vary depending on your environment and the bandwidth to the cloud storage. You may choose the default value.

- 13 Review the configuration summary, and then click **Finish**.

Backup Exec creates a cloud storage device. You must restart Backup Exec services to bring the new device online.

- 14 In the window that prompts you to restart the Backup Exec services, click **Yes**.

After services restart, Backup Exec displays the new cloud storage location in the **All Storage** list.

Microsoft Azure cloud-based storage

The following sections provide information about the requirements to configure a Microsoft Azure cloud-based storage device and configure the storage for Microsoft Azure storage in Backup Exec.

See [“Configuring storage for Microsoft Azure cloud storage”](#) on page 376.

Requirements for configuring a Microsoft Azure cloud-based storage device

Backup Exec cloud connector enables Backup Exec to back up data to and restore data from Microsoft Azure cloud-based storage device.

Review the following requirements before configuring a Microsoft Azure cloud-based storage device:

- You must configure the Microsoft Azure storage account and at least one storage access key (primary access key or secondary access key).
- Ensure that you have already created the blob storage containers for the storage account. Blob storage containers represent a logical unit of storage on the cloud-based storage device.

Note: As a best practice, you should create specific containers to use exclusively with Backup Exec.

Each cloud storage device must use a different container. Do not use the same container for multiple cloud storage devices even if these devices are configured on different Backup Exec servers.

- Backup Exec supports two access tiers of Microsoft Azure. The options are **Hot** (for frequently accessed data) and **Cool** (for infrequently accessed data). For more information about the access tiers, refer to the following link: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>
- Ensure that the container names meet the following Backup Exec requirements:
 - Container names can contain lowercase letters, numbers, and dashes (or hyphens)
 - Containers names cannot begin with a dash (or a hyphen)The containers are not available for use in Backup Exec if the container name does not comply with the container naming convention.
To review the list of supported regions, refer to the Backup Exec hardware compatibility list.

See [“Configuring storage for Microsoft Azure cloud storage”](#) on page 376.

Configuring storage for Microsoft Azure cloud storage

You can configure a cloud-based storage device for the Microsoft Azure cloud storage, and then back up data to it.

Note: For cloud-based storage devices, by default, the **Do not verify data for this job** option is now selected in the **Backup Options**. Cloud vendors charge for operations that read data from and write data to the cloud. To avoid charges for reading data during the verify operation of a backup or duplicate job, this option is selected by default.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

See [“Requirements for configuring a Microsoft Azure cloud-based storage device”](#) on page 375.

To configure storage for Microsoft Azure cloud storage in Backup Exec

- 1 On the **Backup and Restore** tab, click the **Back Up > Backup to Cloud** option and then click **OK**.

Or, on the **Storage** tab, in the **Configure Storage** group, click **Cloud Storage** and then click **Next**.

Or, on the **Storage** tab, click **Configure Cloud Storage**.
- 2 Enter a name and description for the cloud storage device, and then click **Next**.
- 3 From the list of cloud storage providers, select **Azure**, and then click **Next**.
- 4 In the **Cloud storage** field, select the Microsoft Azure server name from the drop-down list.
- 5 In the **Logon account** field, select an account from the drop-down list or click **Add/Edit** to add an account.
- 6 On the **Logon Account Selection** dialog box, click **Add**.
- 7 On the **Add Logon Credentials** dialog box, do the following:
 - In the **User name** field, type the Microsoft Azure storage user name.
 - In the **Password** field, type the Microsoft Azure storage account access key. You can enter the primary access key or the secondary access key.
 - In the **Confirm password** field, type the Microsoft Azure storage account access key again.
 - In the **Account name** field, type a name for this logon account.
The Backup Exec user interface displays this name as the cloud storage device name in all storage device options lists.
- 8 Click **OK** twice.
- 9 Select the Microsoft Azure logon account that you created in step 7, and then click **Next**.
- 10 Select a container from the list of containers that are associated with the server name and then click **Next**.
- 11 Specify how many concurrent operations can run at the same time on this cloud storage device, and then click **Next**.

This setting determines the number of jobs that can run at the same time on this device. The suitable value for this setting may vary depending on your environment and the bandwidth to the cloud storage. You may choose the default value.

- 12 Review the configuration summary, and then click **Finish**.

Backup Exec creates a cloud storage device. You must restart Backup Exec services to bring the new device online.

- 13 In the window that prompts you to restart the Backup Exec services, click **Yes**.

After services restart, Backup Exec displays the new cloud storage location in the **All Storage** list.

Private cloud-based storage

The following sections provide information about the requirements to configure a private cloud-based storage device and configure the storage for private storage in Backup Exec. There are also sections that provide information about how to create, view, edit, and delete a cloud instance for a private cloud.

See [“Requirements for configuring a private cloud-based storage device”](#) on page 378.

See [“Configuring storage for a private cloud-based storage device”](#) on page 380.

See [“Creating a cloud instance for a private cloud”](#) on page 379.

See [“Viewing and editing existing cloud instances for a private cloud”](#) on page 382.

See [“Deleting a cloud instance for a private cloud”](#) on page 382.

Requirements for configuring a private cloud-based storage device

Backup Exec cloud connector enables Backup Exec to back up data to and restore data from a private cloud-based storage device.

Review the following requirements before configuring a private cloud-based storage device:

- Configure the private cloud-based storage server in your environment.
Create an account to access the private cloud-based storage server before you configure a cloud storage device in Backup Exec. You must also copy the private cloud-based storage server's access key ID and secret access key from the corresponding portal.
- Ensure that buckets are already created in the private cloud storage server.
Buckets represent a logical unit of storage on the cloud-based storage device.

Note: As a best practice, you should create specific buckets to use exclusively with Backup Exec.

Each cloud storage device must use a different bucket. Do not use the same bucket for multiple cloud storage devices even if these devices are configured on different Backup Exec servers.

- Ensure that the bucket names meet the following Backup Exec requirements:
 - Bucket names can contain lowercase letters, numbers, and dashes (or hyphens)
 - Bucket names cannot begin with a dash (or a hyphen)

Note: The buckets are not available for use in Backup Exec if the bucket name does not comply with the bucket naming convention.

- Create a cloud instance for the private cloud storage server.
See [“Creating a cloud instance for a private cloud”](#) on page 379.
- Ensure that the private cloud-based storage server has a Certificate Authority (CA)-signed certificate. Backup Exec supports only CA-signed certificates while it communicates with the private cloud storage in the SSL mode. If it does not have the CA-signed certificate, data transfer between Backup Exec and the private cloud provider may fail in the SSL mode.
While creating a cloud instance you can decide to use the SSL protocol. Backup Exec supports backup to a private cloud storage even if you decide not to use the SSL protocol when you create a cloud instance.
To review the list of supported regions, refer to the Backup Exec hardware compatibility list.

See [“Configuring storage for a private cloud-based storage device”](#) on page 380.

Creating a cloud instance for a private cloud

Before configuring a cloud storage device for a private cloud provider, you must create a custom cloud instance for the private cloud storage server.

Type the following command in BEMCLI to create a custom cloud instance:

```
New-BECloudInstance
```

Before you create a storage device for a private cloud, you must run this command.

Ensure that the cloud instance name meets the following Backup Exec requirements:

- It can contain letters, numbers, and dashes (or hyphens).
- It cannot begin with a dash (or a hyphen).

Note: You must create only one cloud instance on a Backup Exec server for a private cloud storage server. You can create another cloud instance on a different Backup Exec server for the same private cloud storage server.

For more information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:<Backup Exec install path>\Backup Exec

See [“Configuring storage for a private cloud-based storage device”](#) on page 380.

See [“Requirements for configuring a private cloud-based storage device”](#) on page 378.

Configuring storage for a private cloud-based storage device

You can configure a private cloud-based storage device, and then back up data to it. Before configuring a cloud storage device for a private cloud storage server, you must create a custom cloud instance for the private cloud storage server.

Note: For cloud-based storage devices, by default, the **Do not verify data for this job** option is now selected in the **Backup Options**. Cloud vendors charge for operations that read data from and write data to the cloud. To avoid charges for reading data during the verify operation of a backup or duplicate job, this option is selected by default.

See [“Configuring automatic verify operations for backup jobs”](#) on page 661.

See [“Creating a cloud instance for a private cloud”](#) on page 379.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Requirements for configuring a private cloud-based storage device”](#) on page 378.

To configure storage for a private cloud-based storage device

- 1 On the **Backup and Restore** tab, click the **Back Up > Backup to Cloud** option and then click **OK**.

Or, on the **Storage** tab, in the **Configure Storage** group, click **Cloud Storage** and then click **Next**.

Or, on the **Storage** tab, click **Configure Cloud Storage**.

- 2 Enter a name and description for the cloud storage device, and then click **Next**.
- 3 From the list of cloud storage providers, select **S3**, and then click **Next**.
- 4 In the **Cloud storage** field, select the private cloud-based storage server name from the drop-down list.
- 5 In the **Logon account** field, select an account from the drop-down list or click **Add/Edit** to add an account.
- 6 On the **Logon Account Selection** dialog box, click **Add**.
- 7 On the **Add Logon Credentials** dialog box, do the following:
 - In the **User name** field, type the private cloud-based storage account access key ID.
 - In the **Password** field, type the private cloud-based storage device secret access key.
 - In the **Confirm password** field, type the private cloud-based storage account secret access key again.
 - In the **Account name** field, type a name for this logon account.

The Backup Exec user interface displays this name as the cloud storage device name in all storage device options lists.

- 8 Click **OK** twice.
- 9 Select the private cloud-based storage logon account that you created in step 7, and then click **Next**.
- 10 Select a bucket from the list of buckets that are associated with the server name and then click **Next**.
- 11 Specify how many concurrent operations can run at the same time on this cloud storage device, and then click **Next**.

This setting determines the number of jobs that can run at the same time on this device. The suitable value for this setting may vary depending on your environment and the bandwidth to the cloud storage. You may choose the default value.

- 12 Review the configuration summary, and then click **Finish**.

Backup Exec creates a cloud storage device. You must restart Backup Exec services to bring the new device online.

- 13 In the window that prompts you to restart the Backup Exec services, click **Yes**.

After services restart, Backup Exec displays the new cloud storage location in the **All Storage** list.

Viewing and editing existing cloud instances for a private cloud

Type the following command in BEMCLI to view a custom cloud instance:

```
Get-BECloudInstance
```

Type the following command in BEMCLI to edit a custom cloud instance:

```
Set-BECloudInstance
```

Note: You cannot edit the cloud instance name and cloud provider of a custom cloud instance.

You must create only one cloud instance on a Backup Exec server for a private cloud storage server. You can create another cloud instance on a different Backup Exec server for the same private cloud storage server.

For more information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:<Backup Exec install path>\Backup Exec

See [“Creating a cloud instance for a private cloud”](#) on page 379.

See [“Deleting a cloud instance for a private cloud”](#) on page 382.

Deleting a cloud instance for a private cloud

Type the following command in BEMCLI to delete a custom cloud instance:

```
Remove-BECloudInstance
```

Note: Before deleting a cloud instance, delete the private cloud storage devices that are attached to that cloud instance.

See [“Deleting a storage device”](#) on page 558.

For more information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:<Backup Exec install path>\Backup Exec

See [“Creating a cloud instance for a private cloud”](#) on page 379.

About S3-Compatible Cloud Storage

Backup Exec provides the S3-compatible cloud storage feature. You can use the S3 provider to configure S3-compatible cloud storage devices with Backup Exec.

When the configuration process is complete, you can create a storage device within the Backup Exec console that can access S3 compatible cloud environments. S3 compatible communications have not been tested in all cloud environments and may not work in some cases.

S3 compatible environments that are not specifically listed in the Backup Exec Hardware Compatibility List are considered alternative configurations. The Backup Exec Hardware Compatibility List defines alternative configurations as:

Alternative Configurations: Our Licensed Software is designed to interoperate with many types of systems, applications, and hardware. Sometimes a customer may choose to use our Licensed Software in an "Alternative Configuration", namely, an environment that has not been validated, approved or verified to operate with our Licensed Software or which does not support such Licensed Software or only supports limited functionality. In most cases, we do not support Alternative Configurations, and we have no obligation to provide Support Services to Licensed Software in an Alternative Configuration. We make no warranty with respect to use of Licensed Software in an Alternative Configuration and any such use is at your own risk. A "Supported Configuration" might be converted into an Alternative Configuration where a vendor modifies one of its components that is part of the original Supported Configuration. As a consequence, your Licensed Software would then be operating in an Alternative Configuration. If you experience a problem with the Licensed Software in an Alternative Configuration or if your issue deals with script that was developed by an unauthorized consulting partner, then we may ask you to reproduce the problem in a Supported Configuration environment. Please note we have no obligation to attempt to resolve problems that cannot be replicated in a Supported Configuration. However, if the problem can be replicated in a Supported Configuration, we will investigate the problem in that Supported Configuration and attempt to resolve it. If the problem cannot be replicated in a Supported Configuration, then we may elect not to work on that problem.

See [“Amazon S3 cloud-based storage”](#) on page 368.

See [“Google cloud-based storage”](#) on page 372.

See [“Private cloud-based storage”](#) on page 378.

See [“Configuring S3-Compatible Cloud Storage with Backup Exec”](#) on page 384.

Configuring S3-Compatible Cloud Storage with Backup Exec

You can configure S3-compatible cloud storage that Backup Exec can access.

To configure S3-Compatible Cloud Storage with Backup Exec

- 1 Configure a cloud instance using the Backup Exec Command Line Interface (BEMCLI) and provide the cloud location and configuration parameters to the Backup Exec server.

To configure a cloud instance, you must pre-configure a user account and buckets in the cloud environment.

For S3, following is an example command:

```
New-BECloudInstance -Name "CloudInstance0001" -Provider  
"compatible-with-s3" - ServiceHost "s3.yourendpoint.com" -UrlStyle  
"Path" -SslMode "Disabled" -HttpPort 80 - HttpsPort 443
```

Where `ServiceHost` is replaced with your cloud server endpoint address.

In this example command, the `SslMode` is `Disabled`. Backup Exec supports the SSL protocol. While creating a cloud instance you can decide to use the SSL protocol. It is recommended that the `SslMode` be `Enabled`.

Ensure that the S3-compatible cloud-based storage server has a Certificate Authority (CA)-signed certificate. Backup Exec supports only CA-signed certificates while it communicates with the S3-compatible cloud storage in the SSL mode. If it does not have the CA-signed certificate, data transfer between Backup Exec and the S3-compatible cloud provider may fail in the SSL mode.

- 2 Create a cloud storage device in Backup Exec using the storage device configuration wizard. Then select the newly created cloud instance and the S3 account credentials that can be used to access the S3-compatible cloud.

See [“Amazon S3 cloud-based storage”](#) on page 368.

See [“Google cloud-based storage”](#) on page 372.

See [“Private cloud-based storage”](#) on page 378.

See [“About S3-Compatible Cloud Storage”](#) on page 383.

Configuring S3-Compatible Cloud Storage with V4 authentication in Backup Exec

You can configure S3-compatible cloud storage with V4 authentication that Backup Exec can access.

To configure S3-Compatible Cloud Storage with V4 authentication in Backup Exec

- 1 Configure a cloud instance using the Backup Exec Command Line Interface (BEMCLI) and provide the cloud location and configuration parameters to the Backup Exec server.

You can also configure the cloud instance using the Generic S3 cloud configurator tool.

To configure a cloud instance, you must pre-configure a user account and buckets in the cloud environment.

For S3-v4, following is an example command:

```
New-BECloudInstance -Name "CloudInstance0002" -Provider  
"compatible-with-s3-v4" - ServiceHost "s3.yourendpoint.com"  
-UrlStyle "Path" -SslMode "Disabled" -HttpPort 80 - HttpsPort 443
```

Where `ServiceHost` is replaced with your cloud server endpoint address.

In this example command, the `SslMode` is `Disabled`. Backup Exec supports the SSL protocol. While creating a cloud instance you can decide to use the SSL protocol. It is recommended that the `SslMode` be `Enabled`.

Ensure that the S3-compatible cloud-based storage server has a Certificate Authority (CA)-signed certificate. Backup Exec supports only CA-signed certificates while it communicates with the S3-compatible cloud storage in the SSL mode. If it does not have the CA-signed certificate, data transfer between Backup Exec and the S3-compatible cloud provider may fail in the SSL mode.

2 Add a region to the cloud instance from BEMCLI.

You can also add this region using the Generic S3 cloud configurator tool.

For S3-v4, following is an example command:

```
New-BECloudRegion -InstanceName "CloudInstance0002"  
-IsDefaultRegion "true" -Name "region name" -Identifier "region  
code" -ServiceHost "s3.yourendpoint.com"
```

In this example command, `-IsDefaultRegion` is set to `true`. This is the default region.

Here is another example command:

```
New-BECloudRegion -InstanceName "CloudInstance0002"  
-IsDefaultRegion "false" -Name "region name" -Identifier "region  
code" -ServiceHost "s3.yourendpoint.com"
```

In this example command, `-IsDefaultRegion` is set to `false`. This is not the default region.

- 3 Run the `Get-BECloudRegion` command to see a list of all available cloud regions.
- 4 Create a cloud storage device in Backup Exec using the storage device configuration wizard. Then select the newly created cloud instance and the S3-v4 account credentials that can be used to access the S3-compatible cloud.

See [“Amazon S3 cloud-based storage”](#) on page 368.
See [“Google cloud-based storage”](#) on page 372.
See [“Private cloud-based storage”](#) on page 378.
See [“About S3-Compatible Cloud Storage”](#) on page 383.

Updating an S3-Compatible Cloud Storage region with V4 authentication

You can update an S3-compatible cloud storage region with V4 authentication that Backup Exec can access. You can only change two parameters, the `ServiceHost` and `IsDefaultRegion`.

To update an S3-Compatible Cloud Storage with V4 authentication

- 1 Configure a cloud instance and a cloud region.
- 2 Update the cloud region parameters.
 - To update the `ServiceHost` parameter of a region, following is an example command:

```
Get-BECloudRegion -Name "region name" | Set-BECloudRegion  
-ServiceHost "s3.yourendpoint.com"
```

- To update the `IsDefaultRegion` parameter of a region, following is an example command:

```
Get-BECloudRegion -Name "region name" | Set-BECloudRegion  
-IsDefaultRegion "true"
```

Removing an S3-Compatible Cloud Storage region with V4 authentication

You can remove an S3-compatible cloud storage region with V4 authentication.

To remove an S3-Compatible Cloud Storage region with V4 authentication

- 1 Configure a cloud instance and a cloud region.
- 2 Remove the cloud region.

Following is an example command:

```
Get-BECloudRegion -Name "region name" | Remove-BECloudRegion
```

- 3 Type `y` to remove the region.

About Backup Exec Cloud Deduplication

Backup Exec supports cloud deduplication where deduplication is enabled on a cloud-based storage device. Along with deduplication to a local disk, you can also perform cloud deduplication using the deduplication storage folder.

The Backup Exec Cloud Deduplication feature helps users to save storage costs and also helps to save network bandwidth.

Backup Exec supports cloud deduplication for the following cloud providers:

- Amazon
- Google
- Microsoft Azure
- Generic S3 providers

Refer to the Backup Exec licensing guide to know the license with which this feature is available.

Backup Exec does not support cloud deduplication for the following cloud provider and storage tiers:

- Provider: Alibaba Cloud

- Glacier (Provider: Amazon)
- Deep Archive (Provider: Amazon)

See [“Notes for Backup Exec Cloud Deduplication”](#) on page 389.

See [“Cloud deduplication storage device”](#) on page 390.

Notes for Backup Exec Cloud Deduplication

The following notes provide information about the Backup Exec Cloud Deduplication feature.

- Cloud deduplication storage devices only support 256-bit AES (PBKDF2) common encryption key.
- Encryption once enabled for cloud deduplication storage devices cannot be disabled but you can change the key. If you did not set an encryption key while creating the cloud deduplication storage device, you can enable it from the device properties page.
- All the data that is sent to the cloud is encrypted. After you enable encryption, you cannot disable it but you can change the encryption key from the properties of the cloud deduplication storage device.
- All cloud deduplication storage devices that are created on a media server use the same encryption key.
- If you create more than one cloud deduplication storage device, you cannot change the encryption key. The new device uses the encryption key that was used to create the first device. In the Backup Exec user interface, the **Encryption key** option is grayed out.
- If you want to use a different encryption key with a cloud deduplication storage device, you can go to the **Storage** tab, select a device, and click **Properties**. In the **Encryption key** option, you can select a new encryption key or add a new key. Restart all Backup Exec services for the key to be used.
- You do not need to change each key of the cloud deduplication storage devices. If you change the key on one of the device properties, it is changed for all the devices that are configured on the media server.
- When you change an encryption key, all keys that the device used must be available on the media server. Backup Exec does not allow you to delete an encryption key based on the following conditions:
 - If the key was used for the cloud deduplication storage devices on the media server.

- If it is an active key that is currently in use by the cloud deduplication storage devices.
- You cannot configure a cloud deduplication storage device on a Managed Backup Exec Server (MBES) from a Central Administration Server (CAS) server. To configure Cloud Deduplication Storage, you need to go to the Managed Backup Exec Server and configure the Cloud Deduplication Storage device.
- When you run a duplicate job to a cloud deduplication storage device, there may be an inconsistency in the byte count of the job, where the byte count in the destination is a little more than the source. Approximately 4 to 6 MB of additional metadata is required per media for the headers. The headers are required for cloud duplication.
- You cannot configure cloud deduplication storage on containers that have immutable storage option or Object Lock enabled.

See [“Cloud deduplication storage device”](#) on page 390.

Cloud deduplication storage device

The following sections provide information about the requirements to configure a cloud deduplication storage device and steps to configure and delete cloud deduplication storage devices from Backup Exec.

See [“Requirements for configuring a cloud deduplication storage device”](#) on page 390.

See [“About cloud immutable \(WORM\) storage support ”](#) on page 393.

See [“Using the cloud admin command-line tool with Amazon S3”](#) on page 394.

See [“Notes for Backup Exec Cloud Deduplication”](#) on page 389.

See [“Configuring a cloud deduplication storage device”](#) on page 398.

See [“Deleting a cloud deduplication storage device”](#) on page 400.

Requirements for configuring a cloud deduplication storage device

Review the following requirements before configuring a cloud deduplication storage device:

- You must first set up a local deduplication storage folder.
- You must set an encryption key on the cloud deduplication storage device that is used for backups to cloud. You can set an encryption key using the Key Management System from Backup Exec.

- You must have sufficient disk cache size to create cloud deduplication disk storage.
See [“Updating disk cache size when you create or import cloud deduplication storage device”](#) on page 391.
- You must use `msdpclutil.exe` to configure backup sets with WORM enabled object lock.
See [“About cloud immutable \(WORM\) storage support”](#) on page 393.

System requirements for configuring cloud deduplication storage

The system requirements are applicable for both Central Admin Server (CAS) and Managed Backup Exec Server (MBES).

There is no limit to the number of cloud deduplication storage devices that you can have for each Backup Exec Server.

Table 9-1 System requirements

Item	Description
Disk Size	Ensure that the minimum free space available is 217 GB per cloud deduplication storage device, on the local deduplication storage.
CPU	Minimum four cores, 2 GB RAM. For example, to protect 64 TB data, your media server needs 8 cores and 8 GB RAM.
RAM	At least 1.5 GB physical memory and 1 dual core processor for each client if you use client side deduplication.

Updating disk cache size when you create or import cloud deduplication storage device

The free disk space available in the local Deduplication Storage volume must be larger than the total of **CloudDataCacheSize**, **CloudMetaCacheSize**, **CloudMapCacheSize**, and **CloudUploadCacheSize** as specified in **contentrouter.cfg**.

You can update the cache size so that the total cache size is less than the free disk space.

To update the cache size

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security > Disk storage lockdown settings > Disable**.
- 3 Provide the System Logon Account credentials along with the reason to disable the lockdown.
- 4 Open the `contentrouter.cfg` file from the following location:
`<location of dedupe storage folder>/etc/puredisk`
- 5 Modify the **CloudDataCacheSize**, **CloudMapCacheSize**, **CloudMetaCacheSize**, and **CloudUploadCacheSize**.

The following table lists the default values for the disk cache size parameters.

CloudDataCacheSize	100 GB
CloudMapCacheSize	5 GB
CloudMetaCacheSize	100 GB
CloudUploadCacheSize	12 GB

Note: Decreasing the cache size from the default values may affect the performance of backup and restore jobs that are performed on the Cloud Deduplication storage device.

It is recommended that the free disk space be greater than the default values of the parameters, that is 217 GB.

For example, your free disk space = 200 GB, CloudDataCacheSize = 100GB, CloudMapCacheSize = 5GB, CloudMetaCacheSize = 100GB, and CloudUploadCacheSize = 12GB. The total cache size is 217 GB. As your free disk space is lesser than your total cache size, the Cloud Deduplication storage device cannot be created or imported. You must modify the parameters to ensure that the total cache size is less than 200 GB. After the cache size is changed (cache size less than free disk space), you can create or import the Cloud Deduplication storage device.

- 6 Save the file and enable the disk storage lockdown setting from **Backup Exec Settings > Network and Security > Disk storage lockdown settings > Enable**.

About cloud immutable (WORM) storage support

Backup Exec now supports Write once, read many (WORM) for cloud storage devices. Backup images that are written to the cloud cannot be written to, overwritten, erased, or reformatted depending on the retention period that is defined. WORM also helps to protect against ransomware and accidental deletion.

The following cloud providers are supported:

- Amazon S3 WORM storage with S3 Object Lock. For more information about Amazon S3 Object Lock, refer to the following link:
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>
- Azure WORM storage with immutable blob storage. For more information about Azure WORM storage, refer to the following link:
<https://docs.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal>

You can use cloud immutable volume management tool to manage immutable cloud volume.

Immutable Cloud Volume (Cloud LSU) is a cloud volume with the following differences than normal cloud volumes:

- The bucket is Object Lock enabled. It is created with the tool `msdpclutil`.
- The bucket policy is attached to the bucket to protect metadata objects of immutable cloud volume.
- A retention range is defined for the cloud volume. The retention of any backup images must be in this range. This range can be defined and changed with `msdpclutil`.
- A specific lifetime is defined for the cloud volume. After the period is completed, the volume expires. You can use `msdpclutil` to extend the lifetime of the cloud volume.
See “Using the cloud admin command-line tool with Amazon S3” on page 394.
See “Using the cloud admin command-line tool with Azure blob storage” on page 395.

Notes for WORM enabled cloud deduplication storage devices

Review the following notes for cloud deduplication storage devices:

- If the storage server supports WORM but the logical storage unit that is used to configure a storage device does not have WORM enabled, the backup sets that are created are not indelible or immutable. To create immutable and indelible backup sets, enable WORM using the storage server and run inventory on the storage device to update the settings.

- If the WORM setting is changed from Disabled to Enabled for the logical storage unit after configuring a storage device, update the backup jobs that are targeted to the storage device, to create immutable and indelible backup sets. If the WORM setting is changed from Enabled to Disabled for the logical storage unit after the configuration of a storage device, backup jobs targeted to the storage device do not create immutable and indelible backup sets.
- When you configure backup jobs, you must select the **Enable Retention Lock** check box. For more information about enabling retention lock, refer to the following topic:
See [“Configuring storage options for backup jobs”](#) on page 652.

Using the cloud admin command-line tool with Amazon S3

The MSDPC cloud admin tool `C:<Backup Exec install path>\msdpclutil.exe` is used to manage immutable cloud volume. A cloud administrator who has the required permissions must run this tool.

You need to use `msdpclutil.exe` to create the bucket with object lock enabled and create a cloud volume in it. If the object lock enabled bucket already exists, you can use `msdpclutil.exe` to create a cloud volume in this bucket.

Before you use this tool, set the following environment variables:

```
set MSDPC_ACCESS_KEY=xxxx
set MSDPC_SECRET_KEY=yyyyyyyyyyyyyy
set MSDPC_REGION=us-east-1
set MSDPC_PROVIDER=amazon
```

For Amazon S3, `MSDPC_ACCESS_KEY` is the AWS access key associated with an IAM user. `MSDPC_SECRET_KEY` is the secret key associated with the access key. `MSDPC_REGION` is the AWS region where the bucket is created or accessed.

To use the cloud admin command-line tool

1 Create a cloud immutable storage volume.

```
C:<Backup Exec install path>\msdpclutil.exe create --bucket
bucketname --volume volumename --mode <mode> --min 1D --max 30D
--live 2021-12-31
```

where, the `<mode>` can be replaced with `GOVERNANCE` or `COMPLIANCE`.

2 List the cloud volumes.

```
C:<Backup Exec install path>\msdpclutil.exe list --bucket
bucketname
```

3 Update the cloud immutable volume min and max retention period.

```
C:<Backup Exec install path>\msdpclutil.exe update range -b  
bucketname -v volumename --min 1D --max 90D
```

The minimum and maximum values are defined by the min and max options. Both values must be between 1 day and 30 years. The maximum value must be less than the volume live duration.

4 Update the cloud immutable volume live duration.

```
C:<Backup Exec install path>\msdpclutil.exe update live -b  
bucketname -v volumename -l 2022-01-31
```

The volume has live period property which is a timestamp. The backup image retention time must be less than this timestamp.

About immutable storage support for Azure blob storage

Backup Exec supports the immutable storage for Azure Blob Storage to store the backup data. For more information about Azure immutable storage, refer to the following topic:

<https://docs.microsoft.com/en-us/azure/storage/blobs/immutable-storage-overview>

You can use one of the following time-based retention policies for immutable blob data:

■ Locked policy

You cannot overwrite or delete the data that is protected using the locked policy for the defined retention period. Once you set a retention period for the data storage, you can extend it but cannot shorten it.

■ Unlocked policy

You cannot overwrite or delete the data that is protected using the locked policy for the defined retention period. Once you set a retention period for the data storage, you can extend, shorten, or delete it.

For information about immutability policies configurations in Azure, refer to the following topic:

<https://docs.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal>

Using the cloud admin command-line tool with Azure blob storage

MSDP cloud admin tool C:<Backup Exec install path>\msdpclutil.exe is used to manage cloud immutable volume. You can create an Azure cloud immutable volume in the following scenarios:

- Azure storage account has enabled version-level immutability support.
- The container is created through Azure portal and has enabled version-level immutability support.
- You use Azure service principal.

To create a cloud volume when version-level immutability support is enabled

1 Set the following environment variables:

```
# export MSDPC_REGION=<your region>
# export MSDPC_PROVIDER=azure
# export MSDPC_ACCESS_KEY=<your storage account>
# export MSDPC_SECRET_KEY=<your access key>
# export MSDPC_ENDPOINT=https://xxxx.blob.core.windows.net/
```

2 Create a cloud immutable volume.

```
C:<Backup Exec install path>\msdpclutil.exe create -b bucketname
-v volumename --mode GOVERNANCE --min 1D --max 30D --live
2022-12-31
```

GOVERNANCE is unlocked policy and COMPLIANCE is locked policy in Azure.

3 List the cloud volumes.

```
C:<Backup Exec install path>\msdpclutil.exe list
```

4 Update the cloud immutable volume min and max retention period.

```
C:<Backup Exec install path>\msdpclutil.exe update range -b
bucketname -v volumename --min 1D --max 90D --unload
```

5 Update the cloud immutable volume live duration.

```
C:<Backup Exec install path>\msdpclutil.exe update live -b
bucketname -v volumename -l 2023-01-31 --unload
```

6 List cloud immutable storage cloud providers.

```
C:<Backup Exec install path>\msdpclutil.exe platform list
```


To create a cloud volume when the container is created through Azure portal and has enabled version-level immutability support

1 Set the following environment variables:

```
# export MSDPC_REGION=<your region>
# export MSDPC_PROVIDER=azure
# export MSDPC_ACCESS_KEY=<your storage account>
# export MSDPC_SECRET_KEY=<your access key>
# export MSDPC_ENDPOINT=https://xxxx.blob.core.windows.net/
```

2 Create a cloud immutable volume.

```
C:<Backup Exec install path>\msdpclutil.exe create -b bucketname
-v volumename --mode GOVERNANCE --min 1D --max 30D --live
2022-12-31
```

GOVERNANCE is unlocked policy and COMPLIANCE is locked policy in Azure.

3 List the cloud volumes.

```
C:<Backup Exec install path>\msdpclutil.exe list
```

4 Update the cloud immutable volume min and max retention period.

```
C:<Backup Exec install path>\msdpclutil.exe update range -b
bucketname -v volumename --min 1D --max 90D --unload
```

5 Update the cloud immutable volume live duration.

```
C:<Backup Exec install path>\msdpclutil.exe update live -b
bucketname -v volumename -l 2023-01-31 --unload
```

6 List cloud immutable storage cloud providers.

```
C:<Backup Exec install path>\msdpclutil.exe platform list
```

To create a cloud volume when you use Azure service principal:

1 Set the following environment variables:

```
# export MSDPC_REGION=<your region>
# export MSDPC_PROVIDER=azure
# export MSDPC_ACCESS_KEY=<your storage account>
# export MSDPC_SECRET_KEY=<your access key>
# export MSDPC_ENDPOINT=https://xxxx.blob.core.windows.net/
# export MSDPC_SUBSCRIPTION_ID=<your subscription id >
# export MSDPC_RESOURCE_GROUP=<resource group storage acct is in>
# export AZURE_TENANT_ID=<azure tenant id>
# export AZURE_CLIENT_ID=<azure client id>
# export AZURE_CLIENT_SECRET=<azure client secret>
```

2 Create a cloud immutable volume.

```
C:<Backup Exec install path>\msdpclutil.exe create -b bucketname
-v volumename --mode GOVERNANCE --min 1D --max 30D --live
2022-12-31
```

GOVERNANCE is unlocked policy and COMPLIANCE is locked policy in Azure.

3 List the cloud volumes.

```
C:<Backup Exec install path>\msdpclutil.exe list
```

4 Update the cloud immutable volume min and max retention period.

```
C:<Backup Exec install path>\msdpclutil.exe update range -b
bucketname -v volumename --min 1D --max 90D --unload
```

5 Update the cloud immutable volume live duration.

```
C:<Backup Exec install path>\msdpclutil.exe update live -b
bucketname -v volumename -l 2023-01-31 --unload
```

6 List cloud immutable storage cloud providers.

```
C:<Backup Exec install path>\msdpclutil.exe platform list
```

Configuring a cloud deduplication storage device

You can configure a cloud deduplication storage device, and then back up data to it.

To configure a cloud deduplication storage device

- 1 On the **Backup and Restore** tab, click the **Back Up > Backup to Cloud** option and then click **OK**.

Or, on the **Storage** tab, in the **Configure Storage** group, click **Cloud Storage** and then click **Next**.

Or, on the **Storage** tab, click **Configure Cloud Storage**.

- 2 Enter a name and description for the cloud storage device.

If you want to use the WORM enabled cloud deduplication storage, you must enter the cloud volume name that you used when you ran `msdpclutil.exe`.

- 3 Select the **Enable deduplication to cloud storage** check box to enable the Backup Exec Cloud Deduplication feature.

The **Enable encryption** check box is selected by default.

You must set an encryption key on the cloud deduplication storage device that is used for backups to cloud. You can add an encryption key using the Key Management System in Backup Exec.

- 4 (Optional) Select **I want to import an existing Cloud Deduplication Storage** check box, in the following scenarios:

- An existing cloud deduplication storage device is deleted.
- During disaster recovery, when you want to recover data for an existing cloud deduplication device.

Note: Ensure that the name of the cloud deduplication device along with the bucket or storage container name matches the existing device.

If Cloud Deduplication Storage is encrypted, first add the KMS keys that are associated with the existing Cloud Deduplication Storage and then perform the import operation.

After the import is completed, you must restart the Backup Exec services for the new device to come online. You must also run the inventory catalog for the backup sets to be available.

- 5 Select a 256-bit AES (PBKDF2) encryption key from the **Encryption key** drop-down list.

The list displays **None** if the key is not selected.

- 6 (Optional) Click **Add Keys** to add an encryption key. In **Encryption type** you must only select **256 bit AES (PBKDF2)**.

See [“Creating encryption keys”](#) on page 734.

- 7 Click **Next**.
- 8 In the **Cloud storage** field, select the cloud service provider name that you want from the drop-down list.

Based on the service provider that you select refer to any of the following sections to continue configuring the cloud deduplication storage device.

See [“Amazon S3 cloud-based storage”](#) on page 368.

See [“Google cloud-based storage”](#) on page 372.

See [“Microsoft Azure cloud-based storage”](#) on page 375.

See [“About S3-Compatible Cloud Storage”](#) on page 383.

Deleting a cloud deduplication storage device

You can delete a cloud deduplication storage device. To delete the cloud deduplication storage device, `storageId` and `CachePath` are required.

To delete cloud deduplication storage

- 1 On the command prompt, enter the following command:

```
<Backup Exec install path>\pddecfg -a listcloudlsu
```

A list of cloud deduplication storage devices is displayed.

- 2 From the Backup Exec user interface, right-click the cloud deduplication storage, and click **Disable**.
- 3 Right-click the cloud deduplication storage and click **Delete**.
- 4 Stop the deduplication service and its monitor service.
- 5 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 6 In the left pane, select **Network and Security > Disk storage lockdown settings > Disable**.
- 7 Delete the storage container or bucket configurations in `spoold` using the following command:

```
<Backup Exec install path>\spoold --removepartition <storageId>
```
- 8 Remove the cache folder using the following command:

```
# rmdir /S <CachePath>
```
- 9 Enable the disk storage lockdown setting from **Backup Exec Settings > Network and Security > Disk storage lockdown settings > Enable**.

- 10 (Optional) Remove the sub-bucket folder from the cloud.
- 11 Start the deduplication service and the monitor service.

Notes for cloud-based storage devices

Glacier and Glacier Deep Archive support

The following information applies to Glacier and Glacier Deep Archive support:

- **Backing up to a cloud storage device:**
When the backup sets are deleted, the storage provider still incurs charges based on their pricing model.
For example: Go to **Backup Exec > Configuration and Settings > Job Defaults > Storage**
Set the **Keep for** value - **1 month**
In Amazon set the time for which you want to keep the backup sets as **3 months**
Even if you delete the data before 3 months, the cloud provider charges based on what is set in Amazon, that is for complete 3 months.
- **Restoring from cloud:**
Restore from Glacier and Deep Archive uses Standard retrieval that can take from 3 to 12 hours depending on the storage device that you selected. The restore operation is time consuming.
- **When you select Glacier or Deep Archive, the **Inventory and Catalog now** and **Catalog** operations are disabled for Glacier and Glacier Deep Archive storage at the device level. You can run the **Inventory** operation independently on this storage device.**
The **Catalog** operation is supported only for an individual media.
You can run the catalog of individual media by performing the following steps in the order listed:
 - Set the registry value from the Backup Exec server to show hidden cloud media.
`HKLM:\SOFTWARE\Veritas\Backup Exec For Windows\BackupExec\User Interface`
DWORD value: **ShowHiddenMedia** set it to **1**.
 - **Catalog individual media by running the following command from BEMCLI:**
`Get-BEMedia OST00000170 | Submit-BECatalogMediaJob`
`OST00000170` is the example media name in this catalog command.

Editing the properties of a cloud-based storage device

You can view all of the properties of a cloud-based storage device and you can change some of the properties.

To edit the properties of a cloud-based storage device

- 1 On the **Storage** tab, double-click the name of the cloud storage device.
- 2 In the left pane, click **Properties**.

3 Change the following properties as needed:

Name	Indicates the user-defined name for this cloud storage device. You can modify this field.
Description	Indicates the user-defined description of this cloud storage device. You can modify this field.
State	Indicates the current state of the device. You cannot modify this field.
Cloud Storage	Indicates the fully qualified name of the server on which the device exists. You cannot modify this field.
Storage location	Indicates the name of the server on which the device exists. You cannot modify this field.
Storage type	Indicates the type of cloud storage device. You cannot modify this field.
Storage tier	Indicates the name of the storage tier. You cannot modify this field. Note: This field displays N/A if the storage tiers are not supported for a cloud storage device.
Bucket/Storage container	Indicates the name of the storage location on the cloud storage device. These storage units are called buckets. You cannot modify this field.
Storage WORM	Indicates whether the Write Once Read Many (WORM) capability is supported for the host server.
Subfolder for Bucket/Storage container	Indicates the name of the subfolder of the storage location on the cloud storage device. These storage units are called buckets. You cannot modify this field.

WORM features	Indicates the WORM features supported by the selected logical storage unit.
WORM indelible minimum interval	Indicates the logical storage unit minimum retention interval specified for the cloud provider.
WORM indelible maximum interval	Indicates the logical storage unit maximum retention interval specified for the cloud provider.
Backup Exec server with device proximity	Indicates the Backup Exec server that has physical or logical proximity to the storage device to run operations for that device. Use this field in a CAS environment.
Host server	Indicates the server to which the local deduplication disk storage belongs.
Cache path	Indicates the location of the local deduplication storage folder.
Logon account	Indicates the name of the logon account that is required to access the device. You can modify this field.
Encryption	Indicates if encryption is enabled for the cloud deduplication storage device. If encryption is enabled, the option is grayed out. You cannot disable the encryption if it already enabled.
Encryption key	Displays the selected encryption key for the cloud deduplication storage device. If you have not enabled encryption, None is displayed. Click Add Keys to add an encryption key. In Encryption type you must select 256-bit AES (PBKDF2). See “Creating encryption keys” on page 734. See “Configuring a cloud deduplication storage device” on page 398.

Concurrent operations	Indicates the maximum number of jobs that you want to run at the same time on this device. You can modify this field.
Client-side deduplication	Indicates whether client-side deduplication is enabled for this device.
Used capacity	Shows the total amount of storage space that is used on this device.
Backup data written	Shows only the amount of the space used for the backed-up data by the cloud or deduplication device.
Deduplication ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.
Connection type	Indicates the type of connection between the Backup Exec server and the cloud storage device. The connection type is Network. You cannot modify this field.
Backup Exec service restart needed	Indicates if the Backup Exec services must be restarted to apply any changes that are made to this device. You cannot modify this field.

- 4
- Click **Apply** to save the changes.

Best practices for using cloud-based storage

Table 9-2 Best practices for using cloud-based storage

Item	Best practice
Data encryption when backing up data to a public cloud storage device	Some public cloud storage providers require encryption to be enabled when backup jobs or duplicate jobs are targeted to cloud storage. All data is secured using SSL during data transfer from Backup Exec to a public cloud-based storage device. However, Backup Exec jobs must have encryption enabled to encrypt the data at-rest in the public cloud storage. See “Using encryption with Backup Exec” on page 729.
Network usage and backups to a cloud storage device	During data transfer, a cloud-based storage device consumes high bandwidth. This might affect the functioning of any other critical applications running at that time. Therefore, it is recommended that you schedule the backup jobs or duplicate backup jobs to a cloud storage device at a time when the bandwidth consumption by other applications is relatively less.

Table 9-2 Best practices for using cloud-based storage (*continued*)

Item	Best practice
Backup Exec CloudConnect Optimizer	Best practices for Backup Exec CloudConnect Optimizer ■ If you have multiple devices for the same cloud provider, you can run the CloudConnect Optimizer job for any one of the devices as the suggested write connection value is applicable for all devices of that cloud server type. ■ Run the CloudConnect Optimizer job close to the backup window so that the Backup Exec CloudConnect Optimizer and backup to cloud jobs run in a similar network environment. Ensure that you have enough time so that the CloudConnect Optimizer job finishes and then the backup to cloud job starts. ■ When you run the CloudConnect Optimizer job, ensure that no backups are running to cloud. This is to ensure that no other backup jobs are sharing the network bandwidth with the Backup Exec CloudConnect Optimizer. ■ It is recommended that you run the CloudConnect Optimizer job, when you configure a cloud storage device, and after you restart the backup exec services. See “About the Backup Exec™ CloudConnect Optimizer” on page 408. See “Configuring the Backup Exec™ CloudConnect Optimizer” on page 409. See “Editing the CloudConnect Optimizer job” on page 410. See “Deleting the CloudConnect Optimizer job” on page 411.

Changing default cloud storage settings

The default cloud storage settings are now enabled in Backup Exec to reduce the storage read write errors that appear when the network is not able to handle the number of connections from the Backup Exec media server to the cloud storage provider.

These connections to cloud are for reading information from the cloud and writing information to the cloud. These are the maximum number of connections that Backup Exec can open.

Currently, the maximum in the connection range is set by default for the read and write connections. It is recommended that you change the values only if there are constant backup failures due to network or bandwidth issues with the set value. In such a case, run the Backup Exec CloudConnect Optimizer in Backup

Exec, which suggests the number of write connections suitable for a backup job in your environment.

The Backup Exec CloudConnect Optimizer displays a job log with the suggested Write Connections value for the specific cloud server type. This value is updated to the Backup Exec settings if you select the **Use the suggested number of Write Connections for the upcoming jobs** check box in the Backup Exec CloudConnect Optimizer.

See [“Configuring the Backup Exec™ CloudConnect Optimizer”](#) on page 409.

You can change the read and write connection values for each type of cloud storage server that is supported by Backup Exec. For backup or restore jobs running on cloud devices, it is recommended that you change the read and write connection values only during network and bandwidth issues. For more information, refer to the Backup Exec Best Practices document.

To change the default storage settings

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Cloud Storage**.
- 3 Edit any of the following options:

Cloud server type	Displays the type of cloud server.
Connection range	Displays the connection range for each type of cloud server.
Read connections	Lets you select the number of read connections for each type of cloud server. The default value for each cloud server is already selected. The default values are set as per the best practices suggested by Backup Exec.
Write connections	Lets you select the number of write connections for each type of cloud server. The default value for each cloud server is already selected. The default values are set as per the best practices suggested by Backup Exec.

About the Backup Exec™ CloudConnect Optimizer

Run the Backup Exec CloudConnect Optimizer to get a suggested Write Connections value that utilizes the available bandwidth. After the CloudConnect

Optimizer job is completed, the log gives you the results of the job and the suggested write connections value.

See [“Changing default cloud storage settings”](#) on page 407.

After you run the CloudConnect Optimizer job for any cloud storage device, the job log displays the suggested Write Connections value for the specific cloud server type. This value is updated to the Backup Exec settings if you select the **Use the suggested number of Write Connections for the upcoming jobs** check box in the CloudConnect Optimizer.

The Write Connections value is specific to a cloud server type. There can be multiple cloud devices with the same cloud server type and the value suggested by the CloudConnect Optimizer would be the same for all the cloud devices of that specific cloud server type. If you select multiple devices that have different cloud providers, you need to need to run the CloudConnect Optimizer job for each device.

It is recommended that you run the CloudConnect Optimizer job close to the backup window and do not run any other backup jobs to the cloud at the same time.

See [“Configuring the Backup Exec™ CloudConnect Optimizer”](#) on page 409.

See [“Editing the CloudConnect Optimizer job”](#) on page 410.

See [“Deleting the CloudConnect Optimizer job”](#) on page 411.

Configuring the Backup Exec™ CloudConnect Optimizer

You can configure and run the CloudConnect Optimizer job to get a suggested Write Connections value.

To configure the CloudConnect Optimizer job

- 1 On the **Storage** tab, right-click a cloud storage device and click **CloudConnect Optimizer**.
- 2 On the **Backup Exec CloudConnect Optimizer** dialog box, in the **General** tab, edit any of the following options.

Job name	Displays a name for the CloudConnect Optimizer job. You can edit the name of the job.
Cloud server type	Displays the type of cloud server.
Use the suggested number of Write Connections for the upcoming jobs	Select the check box to update the suggested value in the Backup Exec settings.

- 3 Click the **Schedule** tab.
- 4 In the **Schedule options**, edit any of the following options.

Run now	Runs the job immediately.
Run on	Schedules the job to run on a specific date and time.

- 5 Click **OK**.

The CloudConnect Optimizer job runs based on whether you have scheduled for it to run immediately or scheduled for a later date.

In case of CAS-MMS scenario, from the Central Admin Server (CAS), you can create a CloudConnect Optimizer job only for the devices that are local to the CAS server or the one shared with the CAS server. For devices that are local to the Managed Media Server (MMS), a CloudConnect Optimizer job can only be created from that particular MMS. If a cloud storage device is shared between CAS and MMS, and you want job to run on MMS, you must to create job from MMS only. The CloudConnect Optimizer job runs on the media server where the job is created.

The CloudConnect Optimizer job is displayed on the **Job Monitor** and **Storage** tab. If the CloudConnect Optimizer job is scheduled for a later date, the job status displays the **Scheduled** status. After the job completes successfully, the **Job Monitor > Job Histories** section and **Storage > Job History** section displays the job details and the job status.

After the CloudConnect Optimizer job completes, the job log displays the suggested Write Connections value. If you have selected the check box to update the suggested value in the Backup Exec settings, this value is updated and used for the upcoming jobs. A confirmation message is displayed that the value is successfully updated in the Backup Exec settings.

See [“Changing default cloud storage settings”](#) on page 407.

See [“About the Backup Exec™ CloudConnect Optimizer”](#) on page 408.

See [“Editing the CloudConnect Optimizer job”](#) on page 410.

See [“Deleting the CloudConnect Optimizer job”](#) on page 411.

Editing the CloudConnect Optimizer job

You can edit a scheduled CloudConnect Optimizer job.

To edit the CloudConnect Optimizer job

- 1 Do one of the following:

To edit a CloudConnect Optimizer job from the Job Monitor tab.	On the Job Monitor tab, right-click the scheduled CloudConnect Optimizer job and then click Edit .

To edit a CloudConnect Optimizer job from the Storage tab.	On the Storage tab, right-click the scheduled CloudConnect Optimizer job and then click Edit .

- 2 On the **CloudConnect Optimizer** dialog box, make any necessary changes.
- 3 Click **OK**.

See [“About the Backup Exec™ CloudConnect Optimizer”](#) on page 408.

See [“Configuring the Backup Exec™ CloudConnect Optimizer”](#) on page 409.

See [“Deleting the CloudConnect Optimizer job”](#) on page 411.

Deleting the CloudConnect Optimizer job

You can delete a scheduled CloudConnect Optimizer job.

To delete the CloudConnect Optimizer job

- 1 Do one of the following:

To delete a CloudConnect Optimizer job from the Job Monitor tab.	On the Job Monitor tab, locate the scheduled CloudConnect Optimizer job that you want to delete.

To delete a CloudConnect Optimizer job from the Storage tab.	On the Storage tab, locate the scheduled CloudConnect Optimizer job that you want to delete.

- 2 Right-click the scheduled job, and then click **Delete**.
- 3 Click **Yes**.

See [“About the Backup Exec™ CloudConnect Optimizer”](#) on page 408.

See [“Configuring the Backup Exec™ CloudConnect Optimizer”](#) on page 409.

See [“Editing the CloudConnect Optimizer job”](#) on page 410.

Generic S3 Configurator

This chapter includes the following topics:

- [About the Generic S3 Configurator](#)
- [Prerequisites for using Generic S3 Configurator](#)
- [Notes for Generic S3 Configurator](#)
- [Creating a cloud instance](#)
- [Deleting a cloud instance](#)
- [Adding a cloud region](#)
- [Viewing the cloud regions](#)
- [Updating a cloud region](#)
- [Deleting a cloud region](#)

About the Generic S3 Configurator

The Generic S3 Configurator utility is used to manage cloud instances configuring S3-compatible cloud storage devices with Backup Exec. The S3-compatible cloud storage is configured using V2 or V4 authentication that Backup Exec can access. This utility can also be used for managing a region.

The Generic S3 Configurator utility is available in English only.

Prerequisites for using Generic S3 Configurator

Review the following prerequisites:

- Ensure that Backup Exec 22.0 or later is installed.

- Ensure that all the Backup Exec services are running.
- Ensure the following:
 - The Service Host points to the correct S3-compatible cloud implementation.
 - The provider name is accurate.
 - The SSL mode is set correctly.

Notes for Generic S3 Configurator

Review the following notes:

- If you use Windows Server 2016 or 2019 as Backup Exec media server, ensure that KB2506143 is installed.
<https://www.microsoft.com/en-in/download/details.aspx?id=34595>
- Backup Exec supports only Certificate Authority (CA)-signed certificates to communicate with cloud storage in the SSL mode.
Ensure that the cloud server has a CA-signed certificate. If it does not have the CA-signed certificate, data transfer between Backup Exec and the cloud provider may fail in the SSL mode. You can disable SSL by setting **SSL Mode** as **Disabled**.
- All the commands can be used in BEMCLI to create an instance or region.

Creating a cloud instance

You can create a new Generic S3 or Generic S3 V4 cloud instance and create a cloud storage.

To create a cloud instance

- 1 Open the S3 Configurator utility.
- 2 Click **Create Cloud Instance**.
- 3 Configure the following preferences:

Name

Enter a name of the new cloud instance. Cloud instance names must meet the following Backup Exec requirements:

- Instance names can contain letters, numbers, and dashes (or hyphens).
- Instance names cannot begin with a dash (or a hyphen).

A cloud instance is not available for use in Backup Exec if the name does not comply with the naming convention.

Service Host	Specify the service host of the cloud instance. A service host must be unique for each cloud instance that is created on the Backup Exec server.
SSL Mode	Specify the SSL mode that Backup Exec uses for communication with the cloud storage server. ■ Full - Use SSL for authentication and data transfer. ■ Disabled - Do not use SSL. ■ AuthenticationOnly - Use SSL for authentication only. Note: It is recommended that you set the SSL Mode to Full or AuthenticationOnly.
URL Style	Specify the URL style of the cloud instance. ■ Path - In a path-style URL, the bucket name is not part of the domain (unless you use a region-specific endpoint). For example: US East (N. Virginia) region endpoint has the following path: http://s3.amazonaws.com/bucket. ■ Virtual - In a virtually hosted URL, the bucket name is part of the domain name in the URL. For example: http://bucket.s3.amazonaws.com ■ Unknown - do not set any value. By default, the URL style is set to Path.
Provider	Specify the provider name of the cloud instance. For S3, the provider value is compatible-with-s3 and for S3V4 authentication, the provider value is compatible-with-s3-v4.

Advanced Menu Change the HTTP or HTTPS port numbers.
The default HTTP and HTTPS values are 80 and 443.

4 Click **Execute Command**.

The BEMCLI output is displayed in the command line. If the command is successful and the cloud instance is created, go to the Backup Exec user interface and create a cloud storage.

```
Name           : cloud_instance_2
Id             : f3f51615-80c5-4cee-b7dc-b12e18726eb6
Provider       : compatible-with-s3-v4
ServiceHost    : s3.au-syd.cloud-object-storage.appdomain.cloud
SslMode        : Disabled
UrlStyle       : Path
HttpPort       : 80
HttpsPort      : 443
```

If the command fails, an error is displayed along with the details.

Deleting a cloud instance

You can delete an existing cloud instance.

To delete a cloud instance

- 1 Click **Delete Cloud Instance**.
- 2 Select the cloud instance you want to remove from the drop-down list and click **Delete**.

The BEMCLI output is displayed in the command line.

Adding a cloud region

You can create a cloud region. Any cloud region that you create must use the Generic S3 v4 version cloud instances.

To create a cloud region

- 1
- Click **Add Region**.
- All existing Generic S3V4 cloud instances are listed in the **Instance Name** list.
- 2
- Configure any of the following preferences:

Instance	Select the Generic S3 V4 instance for which you want to add a region.
Region Name	Enter a name for the region that you want to create.
Identifier	Enter the region's location constraint (such as ID). Select the Set as Default Region check box to make this region your default region. You must select the check box when you create the first region. For all subsequent regions, you can choose to not select the check box.
Service Host	Enter the host name of the relevant cloud region. When you create the first region the host name must match the host name that you used for the cloud instance

- 3
- Click **Execute Command**.
- The BEMCLI output is displayed in the command line. If the command is successful, the region is created.

```
Name      : region_name
Id        : e40b4022-531a-4b15-a010-2c2d12dba5fb
IsDefaultRegion : True
InstanceName : cloud_instance_2
Identifier   : au-syd-standard
ServiceHost  : s3.au-syd.cloud-object-storage.apdomain.cloud
```

Viewing the cloud regions

You can view the existing list of cloud regions with their details.

Click **Get Regions** and the BEMCLI output is displayed in the command line.

Updating a cloud region

You can update an existing cloud region.

To update a cloud region

- 1 Click **Update Region**.

All the existing cloud regions are listed in the **Region Name** list.

- 2 Select the instance for which you want to update the region.
- 3 Select the **Set as Default Region** check box to make this region your default region.
- 4 Select the **Service Host** check box then enter the host name.
- 5 Click **Execute Command**.

The BEMCLI output is displayed in the command line with the updated region details.

```
Name      : region_name
Id        : e40b4022-531a-4b15-a010-2c2d12dba5fb
IsDefaultRegion : False
InstanceName : cloud_instance_2
Identifier   : au-syd-standard
ServiceHost  : s3.au-syd.cloud-object-storage.appdomain.cloud
```

Deleting a cloud region

You can delete an existing cloud region.

To delete a cloud region

- 1 Click **Remove Region**.

All the existing cloud regions are listed in the **Region Name** list.

- 2 Click **Execute Command**.

The BEMCLI output is displayed in the command line.

OpenStorage devices

This chapter includes the following topics:

- [Prerequisites for configuring OpenStorage devices](#)
- [Configuring an OpenStorage device](#)
- [Notes for OpenStorage devices](#)
- [Editing the properties of an OpenStorage device](#)
- [Data Lifecycle Management for WORM enabled OpenStorage devices](#)
- [Specifying a Backup Exec server that has proximity to a shared OpenStorage device](#)

Prerequisites for configuring OpenStorage devices

Before you configure an OpenStorage device, have the following information available:

- What you want to name the OpenStorage device.
- The name of the provider for the OpenStorage device.
- The name of the server and the logon account to use to access the OpenStorage device.
- Whether you want to enable encryption while data is transmitted to the OpenStorage device and while the data is stored on it.
- The number of concurrent operations to run on the device. This setting determines the number of jobs can run at the same time on this device. The number of jobs vary depending on your hardware and environment, so you may need to adjust this setting more than once. It is recommended that you

set the number low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.

- A configured OpenStorage device with the appropriate plug-in installed on the Backup Exec server.
- Whether WORM capability is enabled on the storage server and logical storage unit of an OpenStorage device.
- The minimum and the maximum retention period that is defined as per requirements. This period applies to the backup images that are stored on the logical storage unit.

Configuring an OpenStorage device

OpenStorage is an Arctera technology that allows intelligent disk devices to integrate with Backup Exec.

You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.

Some intelligent disk devices can include multiple logical storage units. However, each logical storage unit is added as a single OpenStorage device. When you add an OpenStorage device, Backup Exec can automatically locate the logical storage units on that device.

Note: When you delete or erase the media from an OpenStorage device, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

If you use Backup Exec Central Admin Server feature, you can share an OpenStorage device between multiple Backup Exec servers. You can enable sharing when you add an OpenStorage device. You can select new Backup Exec servers to share an OpenStorage device or remove the sharing ability for Backup Exec servers at any time. You can specify a different Backup Exec server on which to run backup operations when the OpenStorage device is the source of a job, such as a duplicate job, or a verify job.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 995.

Backup Exec now supports Write once, read many (WORM) on OpenStorage devices. Backup images written to the OpenStorage devices cannot be written to, overwritten, erased, or reformatted depending on the retention period that is defined. WORM also helps to protect against ransomware and accidental deletion.

Before you configure an OpenStorage device, have the following information available:

- What you want to name the OpenStorage device.
- The name of the provider for the OpenStorage device.
- The name of the server and the logon account to use to access the OpenStorage device. You cannot use the System Logon Account. It is recommended that you select or create a logon account that you use exclusively for the deduplication disk storage device. You should not use this account for any other purpose. This account should not contain credentials that are subject to password update policies.
- Whether you want to enable encryption while data is transmitted to the OpenStorage device and while the data is stored on it.
- The number of concurrent operations to run on the device. This setting determines the number of jobs can run at the same time on this device. The number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. It is recommended that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.

To prevent a deduplication job from being sent to a non-deduplication device in a storage pool, you cannot add an OpenStorage device to any storage pools.

To configure an OpenStorage device

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server feature is not installed	Click Network storage , and then click Next .

If the Central Admin Server feature is installed	Do the following in the order listed: ■ Select a Backup Exec server, and then click Next. ■ Click Network storage, and then click Next.

3 Click **OpenStorage**, and then click **Next**.

4 Enter a name and description for the OpenStorage device, and then click **Next**.

5 Do one of the following:

If the provider for the OpenStorage device is listed

Select the provider, and then click **Next**.

If the provider for the OpenStorage device is not listed

Do the following in the order listed:

- Select **My provider is not listed here**, and then click **Next**.
- Enter the provider name, and then click **Next**.

- 6 Enter the name of the server and the logon account to use to access the OpenStorage device, and then click **Next**.
- 7 Select the name of the logical storage unit that you want to use to configure the OpenStorage device, and then click **Next**.
- 8 Specify the number of concurrent operations that can run on the OpenStorage device, and then click **Next**.
- 9 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To create or import the OpenStorage device

Click **Finish**.

See [“Specifying a Backup Exec server that has proximity to a shared OpenStorage device”](#) on page 426.

See [“Prerequisites for configuring OpenStorage devices”](#) on page 418.

See [“Notes for OpenStorage devices”](#) on page 421.

See [“Data Lifecycle Management for WORM enabled OpenStorage devices”](#) on page 425.

Notes for OpenStorage devices

Review the following notes for OpenStorage devices:

- If the storage server supports WORM but the logical storage unit that is used to configure a storage device does not have WORM enabled, the backup sets

that are created are not indelible or immutable. To create immutable and indelible backup sets, enable WORM using the storage server and run inventory on the storage device to update the settings.

- If the WORM setting is changed from Disabled to Enabled for the logical storage unit after configuring a storage device, update the backup jobs that are targeted to the storage device, to create immutable and indelible backup sets. If the WORM setting is changed from Enabled to Disabled for the logical storage unit after configuration of a storage device, backup jobs that are targeted to the storage device do not create immutable and indelible backup sets.

Editing the properties of an OpenStorage device

You can view all of the properties of an OpenStorage device and you can change some of the properties.

To edit the properties for an OpenStorage device

- 1 On the **Storage** tab, double-click the name of the OpenStorage device.
- 2 In the left pane, select **Properties**.

3 Change the following properties as needed:

Name	Indicates the user-defined name for this OpenStorage device.
Description	Indicates the user-defined description of this OpenStorage device.
Backup Exec server with device proximity	Indicates the server that you want to run backup operations when the OpenStorage device is shared between servers in a Central Admin Server feature environment. You can specify a Backup Exec server that has physical or logical proximity to the OpenStorage device. Proximity of the Backup Exec server to the device is an advantage when the device is the source of a job, such as a duplicate job or a verify job, and you want to avoid running the job over a WAN. By default, the Backup Exec server on which you create the OpenStorage device is the server that runs the backup operations for the device.
State	Indicates the current state of the device. You cannot change this property. See “Backup Exec server and storage device states” on page 584.
Host server	Indicates the fully-qualified name of the server on which the device exists.
Server location	Indicates the name or location of the server where the OpenStorage device is located.
Server type	Indicates the type of OpenStorage device, such as PureDisk or the name of the provider for the OpenStorage device.
Storage server WORM	Indicates whether the Write Once Read Many (WORM) capability is supported for the host server.
Logical storage unit	Indicates the logical storage unit used by this OpenStorage device.

Logical storage unit features	Indicates the features supported by the selected logical storage unit.
Logical storage unit indelible minimum interval	Indicates the logical storage unit minimum retention interval specified for the provider.
Logical storage unit indelible maximum interval	Indicates the logical storage unit maximum retention interval specified for the provider.
Logon account	Indicates the name of the logon account that is required to access the device.
Concurrent operations	Indicates the maximum number of jobs that you want to run at the same time on this device.
Split data stream every	Indicates the size at which you want Backup Exec to span to a new image. The default size is 50 GB.
Data stream size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device that is being used.
Stream handler	Indicates whether stream handler is used. Backup Exec sets this option automatically when you select a server type. For some types of devices, this option does not appear at all. If Backup Exec does not set this option, contact the device's vendor for the recommended setting.
Client-side deduplication	Indicates whether client-side deduplication is enabled for this OpenStorage device. Client-side deduplication enables a remote computer to send data directly to an OpenStorage device. By using client-side deduplication, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations.

Disk space to reserve for non-Backup Exec operations

Displays the amount of disk space to set aside for applications other than Backup Exec. The default amount is 5%.

4 Click **Apply** to save the changes.

Data Lifecycle Management for WORM enabled OpenStorage devices

Backup Exec uses data lifecycle management (DLM) to delete expired backup sets and reclaim disk space. For WORM enabled OpenStorage devices there are specific scenarios to be considered.

By default, the **Keep for** value applies to both backup set retention period and image retention lock period.

Table 11-1 DLM scenarios for OpenStorage devices

Scenario	Behavior	Example
Backup set retention period and image retention lock period or the Keep for value are the same.	The catalog and the backup image can be deleted after the period is completed.	Keep for = 2 weeks Retention period on logical storage unit = Minimum - 1 week and maximum - 4 weeks The backup image content and the catalog can be deleted after 2 weeks.
Backup set retention period is changed. A backup set is manually expired (existing period is reduced).	The image can be deleted from the storage server only when the image retention period that is defined on the logical storage unit (at the time that the image was created) has elapsed. The retain reason is displayed on the Backup Exec interface under Storage > Backup Sets.	■ Keep for = 2 weeks ■ Retention period on logical storage unit = Minimum - 1 week Maximum - 4 weeks Backup set expiration period is changed in Backup Exec user interface to 1 week. The image can be deleted from storage and Backup Exec only after 2 weeks.

Table 11-1 DLM scenarios for OpenStorage devices *(continued)*

Scenario	Behavior	Example
Image retention lock period is extended externally (not from Backup Exec).	The image can be deleted from the storage server only when the image retention period that is defined on the logical storage unit (at the time that the image was created) has elapsed. The retain reason is displayed on the Backup Exec interface under Storage > Backup Sets.	■ Keep for = 2 weeks ■ Retention period on logical storage unit = Minimum - 1 week Maximum - 4 weeks Retention lock period interval is changed externally to maximum - 3 weeks. The image can be deleted from storage and Backup Exec after 3 weeks.
Backup set retention period is changed from the Backup Exec user interface.	No effect on image retention.	■ Keep for = 2 weeks ■ Retention period on logical storage unit = Minimum - 1 week Maximum - 4 weeks Backup set retention is changed in Backup Exec user interface to 3 weeks. Backup set is retained for 3 weeks, however the image can be deleted after 2 weeks.

Specifying a Backup Exec server that has proximity to a shared OpenStorage device

You can specify a Backup Exec server that has physical or logical proximity to an OpenStorage device to run operations for that device. Proximity of the Backup Exec server to the device is an advantage when the device is the source of a job, such as a duplicate job or a verify job, and you want to avoid running the job over a WAN.

By default, the Backup Exec server on which you create the device is the server that runs the backup operations for the device. You can specify a Backup Exec server with proximity only if the device is shared between servers in a Central Admin Server feature environment.

If you remove the specified Backup Exec server with device proximity, you are prompted to specify another server.

To specify a Backup Exec server that has proximity to a shared OpenStorage device

- 1 On the **Storage** tab, double-click the name of the device.
- 2 In the left pane, select **Properties**.
- 3 In the **Backup Exec server with device proximity** field, on the drop-down menu, select the server that you want to run operations for the device.
- 4 Click **Apply**.

Microsoft 365

This chapter includes the following topics:

- [About support for Microsoft 365](#)
- [Requirements for Microsoft 365](#)
- [Configuring a tenant for Microsoft 365](#)
- [Backing up Microsoft 365 tenant data](#)
- [Supported workloads and entities for Microsoft 365](#)
- [Restoring Microsoft 365 tenant data](#)
- [Catalog operations for Microsoft 365](#)
- [CAS-MBES scenarios in Microsoft 365](#)
- [Notes for Microsoft 365](#)
- [Limitations of Microsoft 365](#)
- [Recommendations for Microsoft 365](#)
- [Best Practices for Microsoft 365](#)
- [OneDrive plugin: Performance and throttling configuration](#)
- [Exchange plugin: Performance and throttling configuration](#)
- [SharePoint plugin: Performance and throttling configuration](#)
- [Teams plugin: Performance and throttling configuration](#)

About support for Microsoft 365

Microsoft 365 providers recommend that you regularly back up the data. Backup Exec enables you to backup and restore Microsoft 365 data.

Backup Exec supports the following workloads:

- Exchange Online
- OneDrive
- SharePoint Online
- Teams
- Entra ID

High-level workflow

1. Ensure that you have the relevant feature installed for Microsoft 365.
2. Add the Microsoft 365 tenant in Backup Exec.
3. Set up the backup storage.
4. Create and run backup jobs.

Backup Exec uses the Forever Incremental technology for Microsoft 365 backups.

5. Create and run restore jobs.

Requirements for Microsoft 365

Review the following requirements before you back up Microsoft 365 data.

General requirements

- You must install the relevant feature for Microsoft 365.
- You must set up a local deduplication storage disk.

When you set up deduplication disk storage ensure that you consider Microsoft 365 data size requirements and other workload requirements.

It is recommended to set the **Concurrent operations** for a deduplication device to 8 or more. The setting helps to enhance performance of simultaneous runs for Microsoft 365 jobs.
- You must have optimum internet bandwidth as tenant data is hosted in the cloud. Backup and restore requires download and upload of data over an internet connection.

Licensing requirements

Microsoft 365 supports the following Subscription-based licensing editions:

- Instance-based metering - This edition includes the Microsoft 365 feature. It allows you to purchase licenses based on the number of data sources that are protected. An instance equals to 10 Microsoft 365 users.
- Capacity-based metering - This edition includes the Microsoft 365 feature. It allows you to protect total number of front-end terabytes (FETB) of your data including Microsoft 365.

For details on updated licensing, see the [Backup Exec Licensing Guide](#).

Configuring a tenant for Microsoft 365

To backup and restore Microsoft 365 data, a tenant must be added in Backup Exec. You can configure, update, and delete a tenant.

To configure a tenant for Microsoft 365

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Add Microsoft 365 Tenant**.
- 3 In the **Configure Tenant** wizard, select the appropriate region where the tenant exists.
- 4 (Optional) To protect public folders, in **Public Folders Administrator**, specify the User Principal Name (UPN).

The UPN must have administrative access to the entire Public Folders hierarchy starting at the root level. The UPN is validated when the backup job is created, during backup browse of the public folder.

For Exchange Online, use the PowerShell script from the following link to ensure that the specified UPN has Owner-level permissions on the root public folder (IPM_SUBTREE) and its subfolders.

<https://aka.ms/PFPermissionScript>

- 5 Click **Next**.
- 6 Select the workloads that you want to configure for the tenant.

Note: You must select at least one workload.

- 7 Select the check box and click **Next**.

- 8 To configure Microsoft 365 tenant, logon to the tenant with global administrator credentials.

Perform the steps in the following order:

- Click **Copy Code** to copy the device code.
The device code expires after 15 minutes. Ensure that you complete the logon process before the code expires.

Note: If the code expires, click the **Refresh Code** link for a new device code.

- Click the authentication link.
<https://microsoft.com/devicelogin>
- Paste the device code.
- Enter the global administrator credentials and complete the login process.
- After the login process is completed, return to the **Configure Tenant** wizard.

- 9 After the authentication is completed successfully, select the check box and then click **Next**.

Azure Active Directory applications are registered in the Azure Active Directory for the tenant. Global administrator privileges are used to grant permissions to the Azure Active Directory applications to backup and restore Microsoft 365 data.

Backup Exec does not save the global administrator credentials. For backup and restore between Backup Exec and tenant, certificate-based authentication is used.

- 10 Click **Finish**.

The tenant is added. The tenant properties list the type of tenant, tenant ID, and AuthCookie details. Backup Exec raises alerts when the certificates are due for expiration. You must update the tenant configuration to renew the certificates.

For Teams, after you configure the tenant, Channel Posts are backed up using Microsoft Graph Export APIs. These are metered APIs and require additional cost.

To view tenant details

- 1 Right-click the tenant name and then click **Details**.
- 2 In the left pane, select the details that you want to view.

To update the tenant configuration

- 1 Right-click the tenant name and then click **Update**.
- 2 Select the configuration that you want to update.
 - **Update Configuration**
 - **Update Public Folder Administrator**
 - **Delete Public Folder Administrator**
 - **Add Application**
 - **Delete Application**
- 3 If you select **Update Configuration**, you can select any additional workloads for the tenant and then select the check box.

The tenant is configured for the existing and selected workloads.

- 4 If you select **Add Application**, specify the number of additional applications that you want to create.

You can only use the additional applications for SharePoint Online. These applications are only assigned SharePoint API permissions.

Backup Exec displays the maximum number of applications that you can create. The applications help to reduce throttling during backup and restore operations and to improve the backup and restore performance of SharePoint data.

- 5 If you select **Delete Application**, specify the number of applications that you want to remove.

Only the applications added during the add application workflow with SharePoint API are deleted.

Backup Exec displays the number of application that you can remove.

- 6 Click **Next**.
- 7 Perform the same steps to update the tenant using the global administrator credentials.

After the applications are created, you can view the list by going into the tenant properties and click **Fetch details**.

To delete the tenant configuration

- 1 Right-click the tenant name and then click **Remove**.
If you remove the tenant, all the jobs that are associated with the tenant are deleted.
- 2 Click **Yes**.
- 3 Perform steps to delete the application that is registered when you added a new tenant, using the global administrator credentials.

Backing up Microsoft 365 tenant data

You can configure a backup job for Microsoft 365 tenants. For primary backups, the Forever Incremental, Consolidate full, and Custom full templates only support deduplication disk storage. Duplicate backups can be targeted to any supported storage. When you create a forever incremental backup, a full backup is taken followed by incremental backups. A Consolidate Full backup is run by consolidating the previous set of full and incremental backups. Subsequent incremental backups use the consolidate full backup as a baseline to determine changes.

To back up Microsoft 365 servers

- 1 On the **Backup and Restore** tab, select a Microsoft 365 tenant from the list of servers.
- 2 Do one of the following:
 - In the **Backups** group at the top of the screen, click **Backup > Create Microsoft 365 Backup**.
 - On the **Backup and Restore** tab, in the **Groups** pane, right-click the server group that you want to back up and click **Backup > Create Microsoft 365 Backup**.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit** to add or remove resources from the backup selection list.
- 4 On the **Backup Selections** dialog box select the type of data to be backed up.

User

You can select individual users and then select the type of data you want to back up.

- Mailbox
- Archive Mailbox
- OneDrive
- User Site

You can select all users, individual users, or use the search to filter for a specific user.

Group

You can select a group and then select the type of data you want to back up.

- Group Mailbox
- Group Files
- User Mailbox
- User Archive Mailbox
- User OneDrive
- Group Site
- User Site

Depending on the type of data that you select, data for all the users in the selected group are backed up.

You can select all groups, individual groups, or use the search to filter for a specific group.

Public Folders

You cannot select individual Public folders during backup. All the available Public folders are backed up. You can select individual items to restore as the list of public folders is available from the Restore Wizard.

Ensure that the UPN has administrative access to the entire folder hierarchy, because backup runs only for the folders that have access to the hierarchy.

The restore fails if the User Principal Name (UPN) provided during tenant configuration does not have administrative access to the public folders.

Sites

You can select the SharePoint site collections that you want to back up.

You can select all site collections or use the search to filter for a site.

Teams	You can select the teams and then select the type of data that you want to back up. ■ Channel Posts ■ Teams Site and Private Channels Site ■ Teams Data and Settings Teams Data and Settings are always backed up if you select any or both the Channel Posts and Teams Site options. Depending on the type of data that you select, data for the selected teams is backed up. You can select all teams, individual teams, or use the search to filter for teams.
Entra ID	You must select and back up all entities in the Entra ID environment. You cannot select individual Entra ID entities during backup. You can select individual items to restore.

Note: If you select a user and then select a group where the same user is part of the group, the user data is backed up only once.

Note: If you exclude some mailboxes in the **Users** tab for a backup and if those users are part of a selected group in the **Group** tab, then data of the excluded mailboxes is not backed up.

5 On the **Backup Selections** dialog box do any or both of the following:

Select the users or items manually	Select the check boxes corresponding to the users or items for which you want to back up the data.

Click **Apply Filters**

- 1 In **Filter Configuration**, click **Insert**.
- 2 Select the filter, add the keywords, and click **OK**. You can add multiple filters.
- 3 The filters that you added are displayed in **Selection Filters**.

You can select a filter and click **Remove** to delete the filter or click **Clear All** to remove all filters.

The filters with the selected users or items are displayed in **Backup Selections**.

The **Select all** check box is disabled after you use filter option. You cannot apply filters and select all items at the same time.
- 4 Use the sorting buttons to move your selection to the beginning or end of the list.

Even if you use the **Apply Filters** option, any users or items that you select manually always appear on top of the list.

- 6 Click **Selection Summary** to view your backup selection.
- 7 Click **OK**.
- 8 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 9 On the **Backup Options** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.

By default, two templates are available, Forever Incremental and Consolidate Full. You can also add a custom full template.

For more information about Forever Incremental backups, refer to the following information.

See [“Forever Incremental Backup”](#) on page 940.
- 10 On the **Backup Options** dialog box, in the left pane, select **Microsoft 365**, and then select the required parallel streams or workload-specific settings for the backup.

See [“Configuring parallel streams and job settings for Microsoft 365”](#) on page 679.
- 11 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
- 12 Click **OK**.

- 13 (Optional) On the **Backup Definition Properties** dialog box, you can add stages to this backup definition to duplicate your Microsoft 365 data to any supported storage device.
- 14 Click **OK**.

Supported workloads and entities for Microsoft 365

Backup Exec supports the following workloads and plugins:

Table 12-1 Plugins and supported entities

Plugin	Entities
Exchange	User Mailbox, Group Mailbox, Shared Mailbox, Archive Mailbox, Public Folder
OneDrive	User OneDrive, Group Files
SharePoint	User Site, Group Site, Communication Site, Classic Site
Teams	Settings and Channel Tabs, Posts/Conversations, Teams Sites and Private Channels Site
Entra ID	Application registrations, Enterprise applications, Groups, Users

Table 12-2 Entities and descriptions

Entity	Description
OneDrive	
User OneDrive	You can back up user OneDrive for: ■ A user from the User tab. ■ A group that the user belongs to from the Group tab.
Group Files	You can back up group files for a group from the Group tab.
Exchange Online	
User Mailbox	You can back up a user mailbox for: ■ A user from the User tab. ■ A group that the user belongs to from the Group tab.

Table 12-2 Entities and descriptions (*continued*)

Entity	Description
Group Mailbox	You can back up a group mailbox for a group from the Group tab.
Shared Mailbox	A Shared Mailbox is backed up in the same way as a user mailbox. The selection can be the same as a user mailbox.
Archive Mailbox	You can back up an archive mailbox for: ■ A user from the User tab. ■ A group that the user belongs to from the Group tab.
Public folder	You can back up public folders from the Public Folders tab.
SharePoint Online	
User site	You can back up a user site for: ■ A user from the User tab. ■ A group that the user belongs to from the Group tab.
Group site	You can back up a group site for a group from the Group tab. It is applicable to group sites, which are not tied with Teams.
Communication site	You can back up a communication site from the Sites tab.
Classic site	You can back up a classic site from the Sites tab.
Teams	
Channel posts	You can back up channel posts from the Teams tab. Posts from all channels within that Team are backed up.
Teams site	You can back up Microsoft Teams site from the Teams tab.

Table 12-2 Entities and descriptions (*continued*)

Entity	Description
Teams data and settings	You can select the Settings and Channel Tabs to back up metadata for the Teams.
Private Channels Site	You can back up Microsoft Teams Private Channels Site from the Teams tab.
Entra ID	
Enterprise applications	You can back up application configurations and settings associated with enterprise-level applications that are registered and used in the Entra ID environment.
Application registrations	You can back up registration details, settings, and access permissions associated with applications registered in the Entra ID environment.
Groups	You can back up group memberships, access permissions, and settings associated with the groups.
Users	You can back up user profiles, settings, access permissions, and other user-specific details.

Restoring Microsoft 365 tenant data

You can use the Restore Wizard to restore Microsoft 365 tenant data.

If you remove a tenant, before you can restore the data that was backed up using that tenant, you must add the tenant back to the list or catalog the storage device that contains the data that you want to restore.

To restore Microsoft 365 server

- 1 On the **Backup and Restore** tab, select a Microsoft 365 tenant from the list of servers.
- 2 Do one of the following:
 - In the **Restores** group at the top of the screen, click **Restore**.
 - On the **Backup and Restore** tab, in the **Groups** pane, right-click the tenant that you want to restore and click **Restore**.

3 Select any of the following options:

Exchange Online	Exchange Online backups to point-in-time
	Exchange Online backups from backup sets
OneDrive	OneDrive backups to point-in-time
	OneDrive backups from backup sets
SharePoint Online	SharePoint Online backups to point-in-time
	SharePoint Online backups from backup sets
Teams	Teams backups to point-in-time
	Teams backups from backup sets
Entra ID	Entra ID backups to point-in-time
	Entra ID backups from backup sets

4 Select the items that you want to restore.

Show backups from	Displays the beginning and ending dates for the backup sets that you want to include in the selection.
Resource View	Select the Full or the Incremental backup. If you select the point-in-time option, you see a synthesized view of the backup sets, whether incremental or full. You can view the items that were part of the previous backups that are now synthesized and the changes from the most recent backup. An icon indicates the items that are synthesized, and the items that are recently backed up do not have an icon. If you select the backup sets option, you can view the changes from the most recent backup. These items do not have an icon.
Details View	Displays the selected backup sets (selection summary), but you cannot browse or select the contents.

5 Click **Next**.

6 If you select **Exchange Online**, set the following options:

Overwrite option	For an already available item, select if you want to overwrite an item that is already available or if you want to skip that item if it is already available.
Preferred number of parallel streams to use for restore	Specify the number of parallel streams you want to use during restore. By default, the value is set to 1.
Restore to the same Mailbox	Click this option if you want to restore the items to the same mailbox. (Optional) Custom path: Restore to the same mailbox by providing a custom folder inside the destination. Restore to custom subfolders is not supported for Group mailboxes. Re-directed restores are not supported for public folders. You can only restore to the original location.
Restore to a different Mailbox	Click this option if you want to restore the items to a different mailbox. ■ Destination Mailbox Address: Specify the path of the destination mailbox. ■ Destination Mailbox Type: Select the type of the Destination mailbox. Group mailbox can only be restored to a group mailbox. A custom destination path is not supported for restoring to a Group mailbox. When the source and the destination Mailbox types are different, a custom destination path is required. ■ (Optional) Custom path: Specify the custom folder inside the destination. ■ Restore Item Permission: Restores the permissions that were enabled on the item during backup.
Restore to file systems as pst	Click this option if you want to restore the items to a .pst file. You can access the PST file from your Microsoft Outlook. Ensure that you have the 32-bit Outlook version installed on your computer. File path: Browse or enter the target path where you want the .pst file to be created after restore.

Note: Before you start Microsoft 365 restores, you must have a destination mailbox and email addresses created for restore to the same location or redirected restores.

7 If you select **OneDrive**, set the following options:

Overwrite option	For an already available user OneDrive item, select if you want to overwrite an item that is already available or if you want to skip that item if it is already available.
Preferred number of parallel streams to use for restore	Specify the number of parallel streams you want to use during restore. By default, the value is set to 1.
Restore to the same OneDrive	Click this option if you want to restore user OneDrive items to the same location. (Optional) Custom path: Redirect to the same OneDrive but you want to provide a custom folder inside the destination.
Restore to a different OneDrive	Click this option if you want to restore items to a different user OneDrive. ■ Destination User or Group email address: Specify the path of the destination OneDrive. ■ Destination OneDrive Type: Select the type of the destination OneDrive. ■ (Optional) Destination Document Library Name Specify the library to which user OneDrive has to be restored. ■ (Optional) Custom path: Specify the custom folder inside the destination. ■ Restore Item Permission: Restores the permissions that were enabled on the item during backup.
Restore to file system	Enter the target path to restore user OneDrive to a local Backup Exec server. File path For example, the local C or D drive.

8 If you select **SharePoint Online**, set the following options:

Preferred number of parallel streams to use for restore	Specify the number of parallel streams you want to use during restore. By default, the value is set to 1.
Restore only settings, permissions, and structure of unavailable sites/subsites/list	Click this option if you want to restore settings and site structure. Permissions are restored only for the created entities.
Overwrite option	For already available SharePoint site data, select if you want to overwrite the data that is already available or if you want to skip the data if it is already available.
Skip hidden lists and catalogs	Select this option to skip hidden SharePoint Online lists and catalogs for restore. The lists are created by default from SharePoint Online.
Restore Permissions	Select this option to restore permissions of the item and location. Permissions are restored only during disaster recovery scenarios. This check box is selected by default.
Restore to the same SharePoint Site	Click this option if you want to restore data from a SharePoint site to the same location. Version Only the latest version of the document library data in a SharePoint site is restored. Note: For site restore, Backup Exec creates the site if it does not exist. Properties, permissions, and settings for all SharePoint entities are restored along with the data.
Restore to file system	File path Enter the target path to restore data from a SharePoint site to a local Backup Exec server. Version You can restore all versions of the data in the document library. Versioning is only supported for the document library. Or you can just restore the latest version of the document library data from the SharePoint site.

9 If you select **Teams**, set the following options:

Preferred number of parallel streams to use for restore	Specify the number of parallel streams you want to use during restore. By default, the value is set to 1.
Overwrite option for Teams Channel and tabs	For already available Teams channel and tabs, select if you want to overwrite the data that is already available or if you want to skip the data if it is already available.
Select Channel Posts	■ All posts Restore all channel posts for the selected backups. ■ Posts within a date range Restore all channel posts for the selected backup based on the date range. You can also select an individual post for restore.
Restore Teams settings and members	Select the check box to restore all metadata or configuration settings and data related to members for the selected Teams channel. In a disaster recovery scenario, the data is restored even if the check box is not selected.
Skip hidden system lists	Select this option to skip hidden SharePoint Online lists for restore. The lists are created by default from SharePoint Online.
Restore Permissions	Select this option to restore permissions of the item and location. Permissions are restored only during disaster recovery scenarios. This check box is selected by default.
Restore options for channel posts and sites	

Restore to the same Team

■ File version

Restore only the latest version of the SharePoint sites associated with the selected teams.

This option is only applicable to SharePoint sites associated with teams.

■ (Optional) Custom tab name:

Enter the custom tab name for the teams channel. Backup Exec creates the tab name for a channel in a specific format and you have the option to change that.

This option is only applicable to in place restore of channel posts.

Restore to file system

■ File version

You can restore all versions or just restore the latest version of the data from the SharePoint sites associated with teams.

This option is only applicable to SharePoint sites associated with teams.

■ File path

Enter the target path to restore the teams channel posts and sites to a local Backup Exec server.

This option is applicable to channel posts and SharePoint sites associated with teams.

10 If you select **Entra ID**, set the following options:

Preferred number of parallel streams to use for restore

Specify the number of parallel streams you want to use during restore.

By default, the value is set to 1.

Overwrite option

Specify whether to overwrite existing data or skip restore if the data already exists.

Select one of the following options:

■ **Skip if exists:** The data is not restored if it already exists in Entra ID. Restore the data only if it does not exist in Entra ID.

■ **Overwrite:** Overwrite the data if it already exists in Entra ID. Restore the data if it does not exist in Entra ID.

Default password

Enter a password for creating new user objects during the restore job. Ensure that the password complies with the tenant's password guidelines.

Note: This option is enabled only if the **Users** entity is selected during the restore.

Confirm password

Reenter the password.

Restore from recycle bin

Select this option to restore data from the Entra ID recycle bin instead of restoring it from the backup. The deleted data is retained in the Entra ID recycle bin for thirty days.

Force restore of on-premises synchronized directory objects

If you are unable to sync data from on-premises Active Directory to Entra ID, select this option to enforce restoring of data to Entra ID environment.

- 11 Click **Next**.
- 12 Configure any commands that you want to run either before the backup job begins or after the backup job is completed.
See [“Configuring pre/post commands for backup or restore jobs”](#) on page 674.
- 13 Select the name of a recipient to notify when the job is completed. You must have recipients configured to use this option. If the recipient to which you want to send notifications is not in the list, click **Manage Recipients** to add the recipient.
See [“Adding recipients for notification”](#) on page 314.
- 14 Select the **Include the job log in email notifications** check box to send the job log to the recipient that is configured for notification.

15 Specify the job name and select the schedule for the job.

Name	Specify the name of the restore job.
Run now	Runs the job immediately.
Run on	Schedules the job to run on a specific date and time.
Create without a schedule	Creates the job without scheduling it. When you use this option, the job does not run at the time of creation. The job remains unscheduled until you choose to run it.
Reschedule the job if it does not start x hours after its scheduled start	Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.
Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Submit job on hold	Enables you to submit the job, but not run the job until you change the job's hold status.

16 Click **Next**.

17 Review the summary information, and then click **Finish**.

Catalog operations for Microsoft 365

Catalog information is needed to browse the contents of the backup set and perform restore on them. When catalog information about backup sets for Microsoft 365 is stored, the following files are used:

- Image files, which are files containing information about a backup set.
- SQLite database files, which contain detailed item metadata information of backup set. These files are available in a subfolder under the current Catalogs folder.

CAS-MBES scenarios in Microsoft 365

Review the following information that is related to CAS-MBES:

- You can configure Microsoft 365 backup and restore jobs in a CASO-MBES environment. Backup and restore jobs for Microsoft 365 workloads can be

created on the Central Administration Server (CAS) and delegated to run on Managed Backup Exec Servers (MBES).

- When you create a Microsoft 365 workload backup job on CAS, you can choose the server where the job must run. During a restore, the restore job is delegated to the same preferred server that was chosen to run the backup job.
You can view the scheduled backup job on the Job Monitor tab of the CAS. When the job runs, it is delegated to the preferred server and you can monitor the job progress on the Job Monitor tab of the preferred server.
- A backup set that is created by MBES can be browsed from CAS. The detailed item information is fetched from the MBES and displayed on CAS.
- In a CASO environment, for the Forever Incremental, Consolidate Full, or Custom full backup jobs, you must only select a local deduplication storage device hosted on the server where you want to run the backup.
- In a CASO environment, Microsoft 365 always uses distributed catalog mode irrespective of the selected mode during installation.
- In a CASO environment, a backup set that is created by MBES and browsed from CAS does not display detailed item information and cannot be restored from CAS. You must run the restore job from the MBES where the backup job has run.

Notes for Microsoft 365

Note the following points:

General notes

- For primary backups, Forever Incremental, Consolidate full, and Custom full templates only deduplication disk storage is supported. Duplicate backups can be targeted to any supported storage.
- When you run any Microsoft 365 consolidate full job, ensure that no other jobs are running on the deduplication disk storage.
- Support for the Azure China region is discontinued. You can no longer add a new Microsoft 365 tenant for the Azure China region or edit an existing tenant configured for the Azure China region, as it has reached end-of-life.

Exchange notes

- For Exchange Online restore, ensure that the destination mailbox is already created.

- For Exchange Online restore, only the name, comment, permissions, and retention settings for folders are restored.

OneDrive notes

- For OneDrive restore, ensure that the user OneDrive is already created.
- For OneDrive restore, to receive an email invite for direct access permission that is shared with groups, you must enable the **Let people outside the organization email the groups** setting and click **Save**. To enable the setting, go to **Edit group > About > Edit settings**.

Note: This setting is only available in the Microsoft 365 portal and does not exist in Backup Exec.

- For OneDrive, Deep Directory only supports restore of a specific number of items. For more information, refer to the following link:
<https://docs.microsoft.com/en-us/answers/questions/597035/is-there-any-number-restriction-of-sub-folders-in.html>
- For OneDrive Notebook restore, if the Notebooks folder is not available in the root directory, it is created automatically.
- For OneDrive, backup of Preservation Hold Library is supported but you can only restore Preservation Hold Library to a file system.
- For OneDrive, only the latest version of the item at the time of a backup, is protected as part of that backup job.

SharePoint notes

- For SharePoint, if any list item attachments are removed or updated after a backup is completed, those attachments are restored as a part of the restore operation. The restore operation does not delete or modify any attachments, which are not part of the backup data.
- For Sharepoint, TLE (Top Level Entity) and Site name are always the same. The restore browse view also displays the same names.
- If a document is backed up with pending approval status, it is approved during the restore operation.
- After a new SharePoint site is created, Microsoft API may take some time to make the site available for backups in Backup Exec.
- When you restore an event, ensure that you select all the records for repeated events from the restore browse view.

- If you restore a deleted non-repeated event multiple times, duplicate entries of the event are created.
- During file system restore, in the folder created for each event, any character that is not allowed in the file/folder name is replaced with an underscore (_).
- SharePoint Online only restores the latest version of the item and for file system restore, all versions are restored.
- For a list or document, field level unique constraints, validation formula, and custom content type required constraint are not changed during restore. If the constraints and validation formula are changed after backup, it is recommended to keep them consistent with backup.
- The Site Actions within Managed site features are always restored, irrespective of the Backup Exec overwrite settings.
- After upgrade, all items with unique permissions are backed up again. Size of the first incremental can be larger.
- If the restore of settings or permissions fail, the restore job is not marked as failed. The job is marked as success with exception.
- If there is a change in the permission of an item, then the entire item is backed up again.
- Restore of the Restricted View role is supported on entity level. In case of disaster recovery, restricted view role is not restored for any entity.
- If the backup has a default group with new name, then in case of a disaster recovery scenario, you can see two groups. One is the default group created by SharePoint site and the other is created from the backup.

Teams

- Microsoft Teams backup is performed using Graph Export APIs. These are metered APIs and there is a cost associated with the download of each message with these APIs. Refer to the Microsoft documentation for more information about the associated cost.
<https://learn.microsoft.com/en-us/graph/teams-licenses> - model=B
[Get messages across all channels](#) - under model=B
 BackupExec uses model=B based on the recommendation by Microsoft.
- The following applications in channel tabs are protected and can be restored: PowerPoint, Excel, SharePoint, PDF, Word, OneNote, Visio, Power BI, Planner, Document Library, Website, Polly, and Whiteboard. The tabs display applications only if the reference data is available.
 Applications added to a conversation are supported for backup and restore if Teams creates a channel tab for them. Applications that are added directly

to a conversation without creating a tab are not supported for backup and restore. For example, PDF and Power BI.

- After the restore operation, you may see messages with the name of the Azure AD application that was created during tenant configuration. This is because the restore operation works in the context of the Azure AD application and not in the context of any specific user.
- After backup, if a Teams user gets a higher privilege role assigned (for example, owner) then the restore process does not revert this role. The higher privilege role is preserved.
- For file system restore, the folder hierarchy is created based on Teams site url path. Sites include Teams site and Private channels site.

Entra ID notes

- Entra ID objects that are created during restore are created with a new ID. Any applications that reference the old IDs must be updated to use the new IDs.
- In a hybrid environment where synchronization is enabled from on-premises Active Directory to Microsoft Entra ID, restoring directory objects in Microsoft Entra ID involves a three-step process:
 - Restore the relevant directory objects first in the on-premises Active Directory.
 - Wait for the sync from the on-premises Active Directory to Microsoft Entra ID to complete.
 - Restore the relevant directory objects in Microsoft Entra ID to complete the restore of cloud-native attributes.

Limitations of Microsoft 365

The following limitations are observed because of API restrictions from Microsoft 365, Microsoft, and Backup Exec:

Exchange Online

- When you select the **Overwrite** option, from the **Recoverable\Deleted** folder only missing items can be restored. Existing items in the **Recoverable\Deleted** folder cannot be overwritten.
- After restores, folder permissions for some folders (especially calendar folders), are not reflected on the Exchange Online server intermittently.

- When you select the **Overwrite** option, for an item available in a Public Folder, a duplicate is created instead of replacing an existing item.
- When you select the **Retrieve items when Message Read property changes** check box, the item-level changes may not be updated in the immediate incremental backup. The changes are updated in subsequent backups.
- Backup Exec does not back up or restore mailbox-level properties for Exchange Online.
- In Backup Exec, a previously deleted item that was created during restore to Exchange Online is counted as a new mailbox item. If you run the same restore job again, a duplicate of the item is created.
- Backup Exec does not back up the Journal mailbox for Exchange Online.
- For Exchange Online folders, permissions once backed up in the full backup are not backed up again. This is due to an Exchange Online API limitation.
- While restoring data to Exchange Online sometimes an **Access Denied** error is displayed. In this scenario, restore the item to an alternate location.

OneDrive

- Folder size may not be updated if you make changes to the contents of the OneDrive folder, and then take backup immediately.
- OneDrive recycle bin contents are not protected.
- Permissions that are related to links giving access to OneDrive items are not protected.
- Properties of a OneDrive item that is created by a non-OneDrive user are not protected.
- OneNote redirect restore to a document library of the same user or group is not supported.
- In-place restore of OneNote to a non-default document library of a user or group is not supported. In this scenario, you can restore to file system.
- Document library other than the default, which does not have English names, always displays English names. For example, doc1, doc2, and so on.
- If files and folders are shared with others users, then after restore on the shared tab in the other users' account, it shows that files or folders are shared by SharePoint App instead of the name of the user who shared the files.
- Groups that are created from admin center are not displayed in OneDrive and hence in Backup Exec also groups are not displayed.
- OneNote Notebook section sequence is not maintained during restore.

- OneNote Notebook in Azure China region is not protected.
- Direct Access permissions may not be backed up or restored correctly. This is a rare scenario when Graph API does not return correct permissions during backup.
- Sender of the item invite is displayed as SharePoint App user after restoring Direct Access permissions.
- For OneDrive incremental backups, only if the permissions of a folder change, it is not backed up.
- OneDrive restore operation may fail if you have some special characters in the User Principal Name (UPN). You can restore user OneDrive to a local file system, to access the contents.
- OneDrive documents that are protected by a compliance policy and contain sensitive information are skipped during backup without displaying an error. For example, GDPR policy protects personal information, financial information, such as credit card data and account information is protected by a finance policy.

SharePoint Online

- For Teams, Group, and User sites, you need to create the site and then site restore is possible. Properties, permissions, and settings for all entities are restored. But if you create the site, site level properties, permissions and settings are not restored.
- For task list items, tasks and sub-tasks, parent and child relationship is not maintained for already deleted items.
- If there is a network error during the backup, run incremental backups to backup failed or remaining files.
- If you restore a deleted item multiple times, it can create duplicates.
- Backup Exec only restores the latest copies of document library items on SharePoint Online. On the local system, all versions can be restored.
- During restore of SharePoint items, when moderation and approval settings are enabled, the **Modified by**, **Modified date**, **Created date** and **Created by** values are not the same as the backup values. The moderation status is approved after restore. This is due to a Microsoft API limitation.
- If you delete any discussion list after backup, during restore the following is observed:
 - Any parent child relationships are not maintained for the discussion list.
 - Any properties given to the replies in the discussion post are not restored.

- Any hashtags that you created are not restored.
- For a Community Site, discussion lists hashtag (Taxonomy Field) fields are not restored.
- A restore job does not remove or update attachments in existing list items, which are not available during backup.
If the attachments already exist, they are retained.
- SharePoint cannot restore site level lookup columns. This is due to a Microsoft API limitation.
- Any comments given on news items are not backed up or restored by Backup Exec.
- Compliance Policy Center is not supported for backup or restore. This is due to a Microsoft API limitation.
- If files or folders are shared between users, reference of the shared files are backed up the actual file is not backed up. During recovery, the reference is also not restored and you need to share the files again.
- Backup Exec does not back up a file or item does not have a checked in version and a warning message is displayed. Backup Exe does not protect a checked out version of a file or item and no warning message is displayed.
- At the time of backup, if there are checked out and checked in versions of a file or item, only the checked in version is backed up.
- The **Site Collection** tab displays a site by the name, **Team site** that is backed up and restored. The Team site is available by default with Microsoft 365 SharePoint tenant, but the site is not visible in the SharePoint Admin Center.
- File system restore creates a directory structure using the path in the item URL. If the list name does not match with the name given in the URL, the directory name created during file system restore does not match with the list name.

For example: List named "Microfeed" is restored under the directory "Publishfeed".

File System and the folders are restored as:

- "Documents" > "Shared Documents"
- "Composed Looks" > "_catalogs/design"
- "Converted Forms" > "IWConverted Forms"
- "Form Templates > "FormServerTemplates"
- "List Template Gallery" > "_catalogs/It"
- "Master Page Gallery" > "_catalogs/masterpage"

- "Solution Gallery" > "_catalogs/solutions"
- "Theme Gallery" > "_catalogs/theme"
- "User Information List" > "_catalogs/users"
- "Web Part Gallery" > "_catalogs/wp"
- If the read operation of an item fails due to a network error, file download cannot be resumed because of an Microsoft API limitation for older versions. The next backup job attempts to run the failed versions again.
- For a repeated event if any occurrences of the event are deleted or modified, the **Skip if exists** option does not restore the event.
- List lookup columns, List managed metadata columns, Rating settings, and Audience targeting settings are not restored.
- List Views and Forms, Form settings are not backed up and restored.
- System list restore in SharePoint online is not supported.
- In a disaster recovery scenario, Team Site, Group Site and User Site are not created. You must create them manually and then try the restore.
- In Group Calendar Options, the Use this calendar to share members schedule option from Calendar list is not supported.
- The Dialogs property and Automatic Index Management properties are not restored. This is due to a Microsoft API limitation.
- For the **Document Library** and **Picture Library** lists, when you go to **Advanced settings**, the **Site Assets Library** displays **NO** after restore. This is due to a Microsoft API limitation.
- For **Change the look** in site settings, value of the **Current** field is not restored. This is due to a Microsoft API limitation.
- Permissions are restored only if the entity does not exist and is created during restore. If an item exists (site or folder is available), the permissions are not restored irrespective of the overwrite option.
- Teams or Group sites cannot be created during restore so custom groups for the site level cannot be restored if they are deleted. When you restore items that have permissions to those custom groups, the job is marked as success with exception.
- For the **Document Library** list, when you go to **Advanced settings**, the default settings for **Opening Documents in the Browser** are not restored.
- The **Site collection administrators** setting from the **Users and Permissions** site settings is not supported.

- The **About Me** option from the **Permissions Group** setting is not supported.
- Site, subsite, list, column, and page rename for non-English language site and subsite are not supported. This is due to a Microsoft API limitation.
- Description field from List setting is not supported for non-English language site and subsite. This is due to a Microsoft API limitation.
- After restore, the **People and groups** link from **Users and Permissions** site settings are redirected to the **All People** page instead of site members.
- In SharePoint admin center for Teams site, the **Name** field under the site's **Basic info** section is not restored.
- The **List** content type with content type **Post**, which contains the default category column is not supported for in-place restore

Some of the properties that are not supported are listed in the earlier points. There may be other SharePoint properties that are not supported.

Teams

- Backup and restore of group chats and personal chats in Teams is not supported.
- Backup and restore of Shared Channels are not supported.
- Microsoft does not support Teams in the Azure China region. Backup Exec also has no support for Teams in this region.
- If a channel is deleted from Teams, it remains in deleted items for 30 days. A channel with the same name cannot be created for this duration. This is due to a Microsoft API limitation.
If there is a requirement to restore the channel data or settings within the 30 days, you must manually restore the channel from Teams deleted channels and then run the restore job. The job log contains more information.
- There may be an issue while opening the Teams channel conversations HTML file which is restored and attached to a channel tab. This issue is because of the way Microsoft caches the files. Ensure that you have sufficient space on the system drive and try the operation again.
- Some settings in conversations (background image and formatting of the announcement's headline and so on) are not backed up and restored because of the limitations of the Microsoft Graph API.
- Some Teams settings may not get backed up and restored (Team Code, Tags, Hie Team, Show Team, and so on). This is because the graph API used to take backups of Teams settings does not have the provision to retrieve these values.

- Some channel settings (permissions and notifications) may not get backed up and restored because Microsoft Graph API does not have the provision to retrieve the related details.
- Microsoft Graph API cannot retrieve or back up deleted Teams channels.
- After disaster recovery, contents of Wiki tabs are not available. You must click on the restored Wiki tab and then restore its contents separately. This is due to a Microsoft API limitation.
- The OneNote tab links in a Teams channel are not retained in disaster recovery scenarios. The newly created OneNote tab is not configurable as given by Microsoft in the following link:
<https://docs.microsoft.com/en-us/graph/teams-configuring-builtin-tabs>
After the restore is completed, you must create a OneNote tab and associate the restored Notebook with it. You can find the restored notebook on the Teams SharePoint site.
- Restore of archived Teams is only supported in case of disaster recovery. In place restore of archived Teams is not supported.
- In case of disaster recovery, although the documents and files are restored, the custom teams tab cannot access them. You can access the documents or files from the **Files** tab.
- Incremental backups do not capture an updated reaction on a post.
- Restore-browse view for Teams Wiki is not displayed with parent-child hierarchy for Wiki items.
- Teams Wiki page and section order is not maintained after restore.
- The Teams code used for joining a Teams channel directly is not protected. This is due to a Microsoft API limitation.
- The following Teams setting is not protected: **Allow members to upload custom apps**. This is due to a Microsoft limitation.
- If the **Teams** tab is deleted after backup and then the restore operation is run, a new tab is created in each restore session with the same name until next backup is performed.
- In case of disaster recovery, the tab with documents or files displays the following message: **Sorry, this file has been deleted or moved**.
- Teams meeting related data such as the details, recording, and so on are not restored.
- Date range-based backup and restore of Teams posts are run based on created date and last modified date received in the Microsoft Graph API response.

- Deleted post messages are available in the restore browse view, but deleted replies are not available.
- After disaster recovery as the Teams private channel site may or may not be automatically created, the private channel conversation restore may fail. This is a limitation from Microsoft. In a disaster recovery scenario, if restore fails click the files tab of the Private channel to create the channel site and then restore private channel conversation again.
- Embedded images, stickers and code snippets present in Teams posts may not get backed up if the size is greater than 8 MB. This is due to a Backup Exec limitation.
- In case of disaster recovery, the default **Notes** tab in Teams is restored as another tab with the same name. After the restore operation, two **Notes** tabs are displayed. You can access the data from the restored **Notes** tab that was created during disaster recovery.
- If you back up a restored **Posts** tab, after disaster recovery, the earlier restored **Posts** tab is restored successfully, but shows the following message: **Sorry, Something went wrong** when you click the tab. However, you can access the data from the **RestoredPosts** directory on the **Files** tab.

Entra ID

- Restoring user passwords and application secrets is not supported.
- Restoring distribution lists or mail-enabled security groups is not supported.
- Restoring specific object properties is not supported.

Recommendations for Microsoft 365

Note the following recommendations:

- It is recommended that you run a Consolidate Full backup regularly, once a week or once in every two weeks. If you cannot run the backup regularly, it is recommended to run it at least after 30 consecutive incremental backups. Restore operations can then complete in a reasonable amount of time. For jobs that back up large amounts of data, it is recommended to run the Consolidate Full backup after a maximum of 15 consecutive incremental backups. This ensures that the Consolidate Full job can complete in a reasonable amount of time.
- It is recommended that you run a full backup from the source tenant only when required. When you run a full backup from the source tenant, all the data is retrieved from the cloud providers.

- It is recommended that you schedule a consolidate full backup, outside the regular backup window.

- It is recommended to keep the **Verify** option enabled for the forever incremental backup solution.

For simultaneous runs of Microsoft 365 jobs backing up large amounts of data, it is recommended to disable the **Verify** option. Run a separate job to verify the backup set, which helps to complete the jobs within the backup window. If the **Verify** option is enabled as part of backup solution, the next Forever Incremental or Consolidate Full job run is not triggered until the verify operation is completed.

- It is recommended that you run a full backup from the source tenant if a consolidate full job fails.

Best Practices for Microsoft 365

Note the following best practices:

- Ensure that you factor in data increase forecast while creating the deduplication storage device.
- Exclude the deduplication storage location from all anti-virus scanners including Microsoft Defender.
- Set the deduplication storage concurrent operations according to requirements mentioned for Microsoft 365.
- Create a deduplication storage device on Solid-state Drives (SSD) or equivalent storage for enhanced performance.
- Run only one backup job at a time to avoid overlap.
- If the Backup Exec server is a virtual machine, ensure that:
 - CPU reservation is considered.
 - 100% memory is reserved.
- Ensure that the drive in which Backup Exec is installed is different from the drive in which Backup Exec catalog is installed. If your Microsoft 365 jobs are backing up large amounts of data, ensure that the catalog drive has proportionate amount of space.

OneDrive plugin: Performance and throttling configuration

The Backup Exec Microsoft 365 OneDrive plugin is designed to protect a wide variety of data.

A user OneDrive account can contain:

- Small or large files
- Different types of files, such as simple text files, videos, and images
- Large number of files and folders

The plugin ensures that all this data is backed up and restored with minimum throttling and with acceptable performance. All the best practices mentioned by Microsoft are followed by the plugin. For more information, refer to the following link:

<https://docs.microsoft.com/en-us/graph/throttling>

The plugin has built-in scale up and scale down functionality. It helps in scenarios where there are many user OneDrive accounts, which must be backed up and total data size is large (hundreds of GBs to TBs). Although the default tuning works, some tuning may be required for specific scenarios.

Tuning Scenarios

For the following scenarios tuning of the default parameters may be required.

Note: This tuning must be done under the guidance of Technical Support only to avoid issues.

Slow backup speed is observed

If you observe less than expected backup job rate, change some parameters in the `o365_settings.conf` file. If the new values work, you can retain them for subsequent jobs.

Default path

```
C:\Program Files\Veritas\Backup
Exec\SCFPlugins\O365\o365_settings.conf
```

For OneDrive, refer to the section named: **onedrive**

Change the following parameters to improve the backup speed:


```
# Maximum number of requests in a download batch
#backup_max_requests_per_batch = 6
```

Set the value of `backup_max_requests_per_batch` to a higher value. The value ensures that more requests are sent to OneDrive during backup.

```
# Wait time between download batches.
#backup_waittime_between_batches = 5
```

Reduce the value of `backup_waittime_between_batches` to ensure lesser wait time between different batch requests.

Changing these values may increase throttling from Microsoft. While setting these parameters, consider the parallel streams configured in the backup job.

- If you are backing up a single OneDrive account, you can set higher values for these parameters.
- If you are backing up multiple OneDrive accounts in parallel (for example, parallel streams set to 10 or more), the values should be lower.

Throttling errors seen during backup

Throttling may also depend on other factors such as:

- The time during which the backup was run
- Any parallel jobs running at that time

It is recommended to check the throttling behavior for a few runs before changing these parameters.

If high amount of throttling is observed during consecutive backup job runs, tune the following parameters:

```
# Wait time between download batches.
#backup_waittime_between_batches = 5
```

Value of `backup_waittime_between_batches` can be higher so that there is more time between requests and throttling is reduced.

```
# Default number of attempts to retry a failed request.
#backup_max_retry_attempts = 5
```

If multiple requests are failing after retries and this is seen in multiple job runs, then the retry attempts can be higher.

```
# Maximum number of requests in a download batch
#backup_max_requests_per_batch = 6
```

Value of `backup_max_requests_per_batch` can be set to a lower value so that maximum requests per batch are reduced and there is less throttling.

Errors related to properties restore failure seen during restore

When the OneDrive plugin restores an item, the plugin also restores properties and permissions of that item. It requires multiple API calls for a single item, which may result in throttling if there are many items or multiple OneDrives restores running in parallel.

If throttling errors or errors during properties restore are observed, try to restore each OneDrive in separate jobs instead of a single job.

Other options to reduce throttling in restore are to disable properties and permissions restore.

- # Enable/disable properties restore.

By default it is enabled.

```
#restore_in_place_and_redirected_item_properties = 1
```

The property can be used to disable properties restore. It reduces the number of API calls in the item restore and may increase the restore speed. Only item data is restored and no properties (modified by, modified time, and so on) are restored.

- # Enable/disable permissions restore

By default it is enabled.

```
#restore_disable_inplace_item_permissions = 0
```

The property can be used to disable permissions restore. It reduces the number of API calls. If you disable the property, none of the permissions of the item are restored and you must manually assign the permissions. Disable the property if there are multiple throttling errors and no other workaround.

Exchange plugin: Performance and throttling configuration

Performance of backup and restore for M365 Exchange Online depends on the network and the amount of throttling observed from Microsoft 365. M365 Exchange Online throttles requests to maintain optimum performance. Backup Exec also observes the same throttling when sending requests to M365 Exchange Online. Backup Exec follows Microsoft's recommendations to handle throttling.

Considerations before running a backup job

The first full backup can take time. It is important to plan and size your job before starting a backup. Consider factors such as number of mailboxes, type and size of mailboxes, and backup job window before the running the job.

The following recommendations can help increase the speed of backup jobs, particularly the first full job.

Disable EWS throttling

Microsoft 365 has a self-service mechanism by which EWS throttling can be temporarily disabled. It is recommended to disable the throttling temporarily for the first full backup. After the backup is completed, the setting can be enabled.

The Microsoft 365 administrator should open a ticket with Microsoft to disable EWS throttling for a fixed period.

Different types of mailboxes

It is recommended that different types of mailboxes - User Mailboxes, Archive Mailboxes, and Group Mailboxes be backed up in different jobs with different non-overlapping schedules.

Tuning Scenarios

For the following scenarios tuning of the default parameters may be required in the `o365_settings.conf` file. If the new values work, they can be retained for subsequent jobs.

Default path for this file is:

```
C:\Program Files\Veritas\Backup  
Exec\SCFPlugins\O365\o365_settings.conf
```

Note: The changes impact all backup and restore operations on the Backup Exec server. This tuning must be done under the guidance of Technical Support only to avoid issues.

Items failing due to throttling errors during backup

Items throttled during backup are selected in the next backup. Ensure that backups happen when there are no other users accessing the mailboxes. Reducing the number of requests and increasing the number of retries can help in reducing errors.

The following parameters can be reduced:

```
# Number of threads for enumerating folder items during backup.
Microsoft recommended value. (minimum - 2, maximum - 5)
#backup_sync_requests = 5

# Number of threads for backing up items.
(minimum - 2, maximum - 15)
#backup_download_requests = 5

# Number of items to download per request
during backup (minimum - 5, maximum - 100)
backup_items_per_download_request = 80
```

The following parameters can be increased:

```
# Number of retries when requests fail
during backup (minimum - 0, maximum - 10)
#backup_retries_for_requests = 5
```

Items failing due to time out errors during backup

Items can fail with the following error:

End of file or no input: message transfer interrupted or timed out (300 sec max rcv delay) (300 sec max send delay).

Possible reasons for failure:

- Intermittent network errors
- Excessive throttling from Exchange Online

Check for any network issues while connecting with Exchange Online services. Changing the following values can help with time out errors.

The following parameters can be reduced:

```
# Minimum number of items to download per request
to fallback to on network timeout errors
during backup (minimum - 1, maximum - 20)
#backup_min_items_per_download_request = 5
# Maximum total size of all items to download
per request during backup (minimum - 1, maximum - 500)
#backup_max_size_per_download_request_mb = 20
```

The following parameters can be increased:

```
# Network send timeout for requests sent
during backup (minimum - 0, maximum - 300)
#backup_gsoap_send_timeout_secs = 300
# Network receive timeout for requests sent
```

```
during backup (minimum - 0, maximum - 300)
#backup_gsoap_recv_timeout_secs = 300
```

Items failing due to throttling/batch processing errors during restores

For the failed items, you must selectively run the restores for the items to be restored. To avoid throttling failures in any new restores, ensure that restore happens when there are no other users accessing the mailboxes targeted for restore or there is no active backup for those mailboxes.

Reducing the number of requests and increasing the number of retries can help in reducing the errors.

The following parameters can be reduced:

```
# Maximum number of items to upload
in a single request during restore (minimum - 1, maximum - 200)
#restore_max_items_per_upload_batch = 100
# Maximum total size of all items to upload per request
during restore (minimum - 1, maximum - 50)
#restore_max_payload_per_batch_mb = 5
# Number of threads to upload items
during restore (minimum - 1, maximum - 10)
#restore_max_item_upload_tasks = 3
# Maximum number of items to verify in a single request
during restore (minimum - 1, maximum - 200)
#restore_max_items_per_verification_batch = 50
# Number of threads to verify items
during restore (minimum - 1, maximum - 10)
#restore_max_item_verification_tasks = 2
```

The following parameters can be increased:

```
# Number of retries when requests fail
during restore (minimum - 0, maximum - 10)
#restore_max_retries = 5
```

Slow backup when EWS throttling is disabled from Exchange Online

When EWS throttling is disabled from Exchange Online, changing the following settings can increase the speed of backups.

The following parameters can be increased

```
# Number of threads for enumerating folder items during backup.
Microsoft recommended value. (minimum - 2, maximum - 5)
```

```
#backup_sync_requests = 5
# Number of threads for backing up items. (minimum - 2, maximum - 15)
#backup_download_requests = 5
# Number of items to download per request
during backup (minimum - 5, maximum - 100)
backup_items_per_download_request = 80
# Maximum total size of all items to download per request
during backup (minimum - 1, maximum - 500)
#backup_max_size_per_download_request_mb = 20
# Maximum total size of all items (downloaded and across all
outstanding requests) during backup
#backup_max_total_items_download_size_mb = 100
(minimum - 10, maximum - 2000)
```

Note: Changing the settings has no impact when EWS throttling is enabled on Exchange Online and may reduce the job performance. When EWS throttling is enabled on Exchange Online, revert the earlier settings to their original values.

Slow restore when EWS throttling is disabled from Exchange Online

When EWS throttling is disabled from Exchange Online, changing the following settings can increase the speed of restores.

The following parameters can be increased:

```
# Maximum number of items to upload in a
single request during restore (minimum - 1, maximum - 200)
#restore_max_items_per_upload_batch = 100
# Maximum total size of all items to upload per request
during restore (minimum - 1, maximum - 50)
#restore_max_payload_per_batch_mb = 5
# Number of threads to upload items during restore
(minimum - 1, maximum - 10)
#restore_max_item_upload_tasks = 3
# Maximum number of items to verify in a single request
during restore (minimum - 1, maximum - 200)
#restore_max_items_per_verification_batch = 50
# Number of threads to verify items during restore
(minimum - 1, maximum - 10)
#restore_max_item_verification_tasks = 2
```

Note: Changing the settings has no impact when EWS throttling is enabled on Exchange Online and may reduce the job performance. When EWS throttling is enabled on Exchange Online, revert the earlier settings to their original values.

SharePoint plugin: Performance and throttling configuration

BackupExec Microsoft 365 SharePoint Online plugin is designed to protect data.

A SharePoint site can contain:

- Small or large files and folders
- Different types of files (simple text files, videos, and images)
- Lists with different template types (Document libraries, Event lists, System lists)
- Number of subsites

The plugin ensures that all the data is backed up and restored with minimum throttling and with optimum performance.

Performance of backup and restore for SharePoint Online depends on the network and the amount of throttling observed from Microsoft 365. Microsoft 365 SharePoint Online throttle requests to maintain optimum performance. Backup Exec also observes the same throttling when sending requests to Microsoft 365 SharePoint Online.

All the best practices mentioned by Microsoft are followed by the plugin. For more information, refer to the following link:

<https://docs.microsoft.com/en-us/graph/throttling>

<https://learn.microsoft.com/en-us/sharepoint/dev/general-development/how-to-avoid-getting-throttled-or-blocked-in-sharepoint-online>

The plugin has the mechanism to reduce the throttling. It helps in scenarios where there are many SharePoint/User/Group sites, which must be backed up and total data size is large (hundreds of GBs to TBs). Although the default tuning works, some tuning may be required for specific scenarios.

Considerations before running a backup job

The first full backup can take time. It is important to plan and size your job before starting a backup.

Consider factors such as, number of User/Group/SharePoint sites, and backup job window before the running the job.

Tuning Scenarios

For the following scenarios tuning of the default parameters may be required.

Note: This tuning must be done under the guidance of Technical Support only to avoid issues.

Job Tuning Parameters

- Number of threads per stream (site) from the `o365_settings.conf` file.
- Preferred number of streams to use for backup or restore.
- Total SharePoint app to use for the job. Following is the preferred formula:
 Total number of SharePoint applications = Number of threads per stream (site)
 x Number of streams configured for the job.

If you observe less than expected backup job rate, change preferred number of streams or some parameters in the `o365_settings.conf` file. If the new values work, you can retain them for subsequent jobs.

Changing these values may increase throttling from Microsoft. While setting these parameters, consider the parallel streams configured in the backup job. If you are backing up a single site, you can set higher values for these parameters. If you are backing up multiple sites in parallel (for example, parallel streams set to 10 or more), the values should be lower.

Default path:

C:\Program Files\Veritas\Backup Exec\SCFPlugins\O365\o365_settings.conf

For SharePoint, refer to the section named: **[sharepoint]**

Items failing due to throttling errors during backup

Possible reasons for failure:

- Intermittent network errors
- Excessive throttling from SharePoint Online

Items throttled during backup are selected in the next backup. Reducing the number of requests and increasing the number of retries can help in reducing errors.

The following parameters can be reduced:


```
# Number of threads per stream (minimum - 1, maximum - 50)
# backup_num_threads = 3
```

The following parameters can be increased:

```
# backoff time when requests fail during backup
(minimum - 10, maximum - 60)
#backup_backoff_secs = 30
```

```
# Number of retries when requests fail during backup
(minimum - 0, maximum - 20)
#backup_retry_count = 5
```

Items failing due to throttling errors during restore

Possible reasons for failure:

- Intermittent network errors
- Excessive throttling from SharePoint Online

Reducing the number of requests and increasing the number of retries can help in reducing errors.

The following parameters can be reduced:

```
# Number of threads per stream (minimum - 1, maximum - 25)
# restore_num_threads = 5
```

The following parameters can be increased:

```
# backoff time when requests fail during restore
(minimum - 1, maximum - 60)
#restore_backoff_secs = 30
```

```
# Number of retries when requests fail during restore
(minimum - 0, maximum - 20)
#restore_retry_count = 5
```

Teams plugin: Performance and throttling configuration

The Backup Exec Teams plugin is designed to protect a wide variety of data.

Teams can protect:

- Teams posts
- Teams data and settings

- Back-end Teams site content

While protecting this data, the plugin follows all the best practices recommended by Microsoft. This ensures that there is minimum throttling and failures during backup and restore.

<https://learn.microsoft.com/en-us/graph/throttling>

<https://learn.microsoft.com/en-us/graph/throttling-limits>

Tuning scenario

For the following scenario, tuning of default parameters may be required.

Throttling errors seen during backup of Teams posts

Throttling may also depend on other factors such as:

- The time during which the backup was run
- Any parallel jobs running at that time

It is recommended to check the throttling behavior for a few runs before changing the parameter. If high amount of throttling is observed during consecutive backup job runs, tune the following parameter in the `o365_settings.conf` file. If the new values work, you can retain them for subsequent jobs.

Default path:

C:\Program Files\Veritas\Backup Exec\SCFPlugins\O365\o365_settings.conf

For Teams, refer to the section named: **[teams]**

Change the following parameter to improve the success rate:

```
#backup_graph_link_download_batch_size = 20
```

You can reduce the value of this parameter and uncomment the line in the conf file by removing # at the beginning.

Tape storage

This chapter includes the following topics:

- [Support for tape drives and robotic libraries](#)
- [Adding or replacing devices with the Hot-swappable Device Wizard](#)
- [Editing tape drive properties](#)
- [Viewing tape drive statistics](#)
- [Robotic libraries in Backup Exec](#)
- [Creating robotic library partitions](#)
- [Managing tapes](#)
- [Labeling tape media](#)
- [How WORM media is used in Backup Exec](#)
- [Default media vaults](#)
- [Retiring damaged tape media](#)
- [Deleting tape media](#)
- [Erasing tape or disk cartridge media](#)
- [About cataloging tape or disk cartridge media that contains encrypted backup sets](#)
- [Associating tape media with a media set](#)
- [Editing tape media properties](#)
- [Tape media rotation strategies](#)

Support for tape drives and robotic libraries

When you install Backup Exec, it automatically recognizes all tape storage devices that are attached to the Backup Exec server.

Support for tape drives and libraries varies across Backup Exec editions. Please refer to the licensing information specific to your edition for more information.

You can use the **Configure Storage** wizard to perform the following actions for tape storage:

- Partition robotic library slots.
- Install tape device drivers.

Note: You do not need to install tape device drivers if Backup Exec runs on Windows Server 2012 and later. Kernel-mode drivers and tapeinste.exe are no longer installed if Backup Exec runs on Windows Server 2012 and later.

- Replace or add hot-swappable storage on a Backup Exec server without having to restart the server.
- Create media sets to manage the backup data on tapes.

See [“Default media sets”](#) on page 491.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Adding or replacing devices with the Hot-swappable Device Wizard

Use the **Hot-swappable Device Wizard** to replace or add hot-swappable storage on a Backup Exec server without having to restart the server.

If you remove and then reconnect Universal Serial Bus (USB) tape devices to the USB port, you must run the **Hot-swappable Device Wizard** to allow Backup Exec to rediscover the devices.

For iSCSI-attached devices, you must list the device as a **Persistent Target** in the iSCSI control panel applet, and then run the **Hot-swappable Device Wizard**. Listing the device as a **Persistent Target** lets Backup Exec rediscover the device whenever you restart the Backup Exec server.

After you start the **Hot-swappable Device Wizard**, you are prompted to close the Backup Exec Administration Console. The **Hot-swappable Device Wizard**

waits until any jobs that were processing are completed. The wizard pauses the Backup Exec server and stops the Backup Exec services. You can then add or replace any storage devices. The wizard detects the new device or replaced device, and adds information about the device to the Backup Exec database. The wizard is then completed, and you can reopen the Backup Exec Administration Console.

Any new storage appears in the **Storage** tab, and usage statistics for the storage begin accumulating. You can enable the new storage in a storage device pool.

Any replaced storage appears in the **Storage** tab, in the **All Storage** view with a status of Offline.

Note: Start the **Hot-swappable Device Wizard** before you add or replace storage.

To add or replace devices with the Hot-swappable Device Wizard

- 1 Do one of the following:

For iSCSI-attached storage:

In the iSCSI control panel applet, add the storage to the **Persistent Targets** list.

Continue with the next step.

For any other hot-swappable storage:

Continue with the next step.

- 2 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 3 When you are prompted for the type of storage that you want to configure, select **Tape Storage**, and then click **Next**.
- 4 Select **Run the Hot-swappable Device Wizard**, click **Next**, and then follow the on-screen prompts.

Editing tape drive properties

You can edit the following tape drive properties.

See [“Support for tape drives and robotic libraries”](#) on page 472.

To edit tape drive properties

- 1 On the **Storage** tab, double-click the tape drive for which you want to edit properties.
- 2 In the left pane, click **Properties**.

3 Edit any of the following options:

Name	Displays the name of the tape drive. You can edit this field.
Description	Displays the description of the tape drive. You can edit this field.
Hardware compression	Indicates if hardware compression is enabled. If this option is available, this drive is capable of supporting hardware compression. If you configure a job to use hardware compression, and hardware compression is disabled on the device, then hardware compression is unavailable and is not used.

Block size

Displays the size of the blocks of data that are written to new media in this tape drive. The default is the preferred block size.

Some devices (for example, LTO devices) provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a tape drive that supports larger block sizes, you can change the tape drive's block size. However, if the tape drive does not allow a block size as large as you want, reconfigure the host bus adapter or the tape drive. After you reconfigure the hardware and restart the Backup Exec services, check if the block size that you want to use is available.

See the tape drive manufacturer's documentation for help in configuring the device.

Backup Exec does not ensure that the requested block size is supported by that tape drive. You should check the tape drive specifications to make sure that the block size is supported. If the tape drive does not support a block size, it defaults to its standard block size.

If the tape drive does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of data that is sent to the tape drive on each read or write request. The buffer size must be equal to the block size or an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve tape drive performance. Each type of tape drive requires a different buffer size to achieve maximum throughput.

Buffer count

Displays the number of buffers that are allocated for this tape drive.

Depending on the amount of memory in your system, increasing this value may improve device performance. Each type of tape drive requires a different number of buffers to achieve maximum throughput.

If you change the buffer count, you may need to adjust the high water count accordingly.

High water count

Displays the number of buffers to be filled before data is first sent to the tape drive, and any time that the tape drive underruns.

The high water count cannot exceed the buffer count. A value of 0 disables the use of high water logic; that is, each buffer is sent to the device as it is filled.

The default setting provides satisfactory performance in most instances; in some configurations, throughput performance may be increased when other values are specified in this field. If you increase or decrease the buffer count, the high water count should be adjusted accordingly. If a tape drive has a high water count default of 0, it should be left at 0.

Reset to default settings

Returns all of the preferred configuration settings to their defaults.

Read single block mode

Indicates if this tape drive reads only one block of data at a time, regardless of the size of the buffer block.

This option is disabled by default.

Write single block mode

Indicates if this tape drive writes only one block of data at a time. This option provides greater control over the handling of data write errors.

It is recommended that you select this option if the tape drive is shared.

This option is enabled by default.

Read SCSI pass-through mode

Indicates if this tape drive reads data without going through a Microsoft tape device API. This option allows the data to pass directly through the tape drive and allows more detailed information if device errors occur.

This option is disabled by default.

Write SCSI pass-through mode

Indicates if this tape drive writes data without going through the Microsoft tape device API. This option allows data to pass directly through the device driver and allows more detailed information if device errors occur.

It is recommended that you select this option if the tape drive is shared.

This option is enabled by default.

Servers that share this device

Displays the servers that can also use this device.

See [“Sharing storage devices”](#) on page 557.

Media Type

Indicates if barcode rules are enabled for the robotic library that this tape drive is attached to. If barcode rules are enabled, the media types are listed that this tape drive can read from and write to. Backup Exec uses the barcode rules to identify which type of media to use in a drive.

You can configure barcode rules, and enable or disable barcode rules for a robotic library.

See [“Configuring barcode rules for a robotic library ”](#) on page 480.

Can Read From

Indicates if this tape drive can read from the media type.

See [“Configuring barcode rules for a robotic library ”](#) on page 480.

Can Write To

Indicates if this tape drive can write to this media type.

See [“Configuring barcode rules for a robotic library ”](#) on page 480.

- 4 Click **Apply**.

Viewing tape drive statistics

You can view statistics about tape drives.

See [“Support for tape drives and robotic libraries”](#) on page 472.

To view tape drive statistics

- 1 On the **Storage** tab, double-click the tape drive for which you want to view statistics.
- 2 In the left pane, click **Statistics**.

Robotic libraries in Backup Exec

Backup Exec’s Advanced Device and Media Management (ADAMM) feature solves the problems that are associated with typical robotic library modules. Backup Exec accesses all of the media in the robotic library and uses the media that belongs to the specified media set. If the backup job exceeds the capacity of a media, Backup Exec searches all of the media that is contained in the robotic library and finds a suitable media to use.

For example, an operator has a robotic library with six slots. The operator inserts six blank tapes and targets backup jobs to various media sets within the robotic library. Backup Exec automatically allocates available tapes in the robotic library. If a job exceeds the capacity of one tape and another overwritable tape is available in the robotic library, the job automatically continues on that tape. When Backup Exec runs out of tapes, it prompts the operator to import overwritable media.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If more than one media that meets the requirements is found, Backup

Exec selects the media in the lowest-numbered slot. For example, Backup Exec selects media in slot 2 before it selects equivalent media in slot 4.

For restore jobs that use robotic libraries, Backup Exec accesses the source media regardless of its sequential placement in the magazine. For example, if the data for a restore job resides on two media in the magazine, the media do not have to be placed in adjacent slots for Backup Exec to restore the data. If Backup Exec does not find the media that is required for the restore job in the robotic library, an alert is generated that requests the media that is necessary to complete the job.

See [“Requirements for setting up robotic library hardware”](#) on page 479.

See [“Creating robotic library partitions”](#) on page 486.

Requirements for setting up robotic library hardware

You can configure Backup Exec to work with robotic library drives by making associations between the robotic library's drives, robotic arm, and Backup Exec. Backup Exec supports serialized drives. Manual configuration of serialized drives is not required.

You can find a list of supported types of storage in the Backup Exec Hardware Compatibility List.

Ensure that the robotic library hardware is configured as follows:

- Ensure that the robotic arm is set to Random mode. Refer to your robotic library documentation for more information.
- Ensure the following for a multi-LUN robotic library:
 - The controller card is set to support multiple LUNs (if supported).
 - The target LUN for the tape drive is lower than the target LUN for the robotic library.
- Determine which drive is the first drive in the robotic library, and then arrange the SCSI IDs to match the sequence of the drive element addresses. Refer to your robotic library documentation to determine the drive element address for each storage device.
- Ensure that the SCSI ID of the robotic arm precedes the SCSI IDs of the drives in the robotic library. Do not use 0 or 1 because these SCSI IDs are typically reserved for boot devices.

In the following example, if your robotic library has two drives, the drive with the lowest drive element address should be assigned the lower SCSI ID.

Table 13-1 Example configuration for a multi-drive robotic library

Data transfer element (storage devices)	SCSI ID	Drive element address
Robotic arm	4	N/A
Storage device 0	5	00008000
Storage device 1	6	00008001

See [“Robotic libraries in Backup Exec”](#) on page 478.

Inventorying robotic libraries when Backup Exec services start

You can set a default so that all robotic libraries are included in the inventory job whenever Backup Exec services are started. It is recommended that you enable this default if media is often moved between robotic libraries. Backup Exec may take longer to start.

To inventory robotic libraries when Backup Exec services start

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Storage**.
- 3 Click **Inventory robotic libraries when Backup Exec services start**.
- 4 Click **OK**.

See [“Inventorying a storage device”](#) on page 562.

Configuring barcode rules for a robotic library

If you have barcode support for a robotic library that uses different types of drives, you can create a barcode rule. Backup Exec uses the barcode rules to identify which type of media to use in a drive. When Backup Exec reads the barcode rule, it locates the type of media that corresponds to the prefix or suffix. Backup Exec then mounts the media into a drive that accepts that type of media.

The barcode rules apply to all of the robotic libraries that you enable for barcode rules. By default, barcode rules are disabled.

You can use the **Configure Storage** wizard to configure barcode rules, or you can add, edit, and delete barcode rules in the Backup Exec global settings. You can also enable or disable barcode rules in the robotic library's properties.

See [the section called “Configuring barcode rules for a robotic library by using the **Configure Storage** wizard”](#) on page 481.

See [the section called “Adding, editing, and deleting barcode rules by using the Backup Exec global settings”](#) on page 483.

See [the section called “Enabling or disabling barcode rules by using robotic library properties”](#) on page 482.

Configuring barcode rules for a robotic library by using the **Configure Storage** wizard

To configure barcode rules for a robotic library

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server feature is not installed Click **Tape storage**, and then click **Next**.

If the Central Admin Server feature is installed

Do the following in the order listed:

- Select the Backup Exec server on which you want to configure storage, and then click **Next**.
- Click **Tape storage**, and then click **Next**.

3 Click **Configure barcode rules for a robotic library**, and then click **Next**.

4 Select the robotic library that you want to configure barcode rules for, and then click **Next**.

5 To change the available barcode rules in the list, click any of the following:

To add a new barcode rule

Do the following in the order listed:

- Click **New**.
- Click the drop-down menu, and then select a media type
- Type the vendor name, and the barcode prefix and/or suffix.
- Click **OK**.

To edit an existing barcode rule

Do the following in the order listed:

- Select the barcode rule that you want to edit, and then click **Edit**.
- Make any changes, and then click **OK**.

To delete a barcode rule

Select a barcode rule, click **Delete**, and then click **Yes** to confirm the deletion.

- 6 Click **Next**.
- 7 Select a tape drive, and then for any of the listed media types, select a check box to specify if the tape drive can read from and write to that media type.
Click **Next**.
- 8 Review the summary, and then do one of the following:

To change the barcode configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next**.
- Click **Finish**.

To complete the barcode configuration

Click **Finish**.

- 9 Ensure that barcode rules are enabled for the appropriate robotic libraries.

Enabling or disabling barcode rules by using robotic library properties

You can enable or disable barcode rules for a robotic library. By default, barcode rules are disabled for robotic libraries.

To enable or disable barcode rules by using robotic library properties

- 1 On the **Storage** tab, double-click the robotic library for which you want to enable barcode rules.
- 2 In the left pane, click **Properties**.
- 3 In the **Barcode rules** field, in the drop-down menu, do one of the following:

To enable all barcode rules for this robotic library Click **Enabled**.

To disable all barcode rules for this robotic library Click **Disabled**.

- 4 Click **OK**.

Adding, editing, and deleting barcode rules by using the Backup Exec global settings

You can add, edit, or delete the barcode rules for a robotic library in the global settings for Backup Exec. All barcode rules changes apply to all of the robotic libraries for which barcode rules are enabled.

To add, edit, or delete a barcode rule by using the Backup Exec global settings

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Bar Code Rules**.
- 3 Do one of the following:

To add new barcode rule

Click **New**, and then continue with the next step.

To edit an existing barcode rule

Select the barcode rule that you want to edit, click **Edit**, and then continue with the next step.

To delete an existing barcode rule

Do the following in the order listed:

- Select the barcode rule that you want to delete, and then click **Delete**.
- Click **Yes** to confirm the deletion, and then click **OK**.

The barcode rule is deleted.

4 Add or change any of the following information:

Media type	Click the drop-down menu and select or change the media type.
Vendor	Enter or change the name of the robotic library manufacturer. A best practice is to use a vendor name only if there is a specific need for it and you have multiple brands of tape drives that are available This field is limited to 16 characters.
Barcode prefix	Enter or change the prefix of a barcode to which you want to apply this barcode rule. A standard practice is to use a prefix to indicate a media-handling policy, such as CLN for cleaning media, or I for internal use, or O for cartridges that are to be taken offsite. This field is limited to 16 characters.
Barcode suffix	Enter or change the suffix of a barcode to which you want to apply this barcode rule. A standard practice is to use a suffix to indicate the generation of the media, such as L1, L2, and so on. This field is limited to 16 characters.

5 Click **OK**.

Initializing a robotic library when the Backup Exec service starts

You can initialize a robotic library whenever the Backup Exec services start.

During startup, if media is in the drives in the robotic library, Backup Exec attempts to return the media to its original drive. If the media cannot be returned to the drive, an error message appears that prompts you to eject the media from the drive.

You can also create a job to initialize a robotic library.

See [“Initializing a robotic library”](#) on page 564.

See [“Robotic libraries in Backup Exec”](#) on page 478.

To initialize a robotic library when the Backup Exec services start

- 1 On the **Storage** tab, double-click the robotic library that you want to initialize.
- 2 In the left pane, click **Properties**.
- 3 In the **Startup initialization** field, in the drop-down menu, click **Enabled**.
- 4 Click **Apply**.

Defining a cleaning slot

Before submitting a cleaning job, you must define a cleaning slot that contains the cleaning tape.

Make sure that the cleaning tape is located in the slot that you defined as the cleaning slot. After defining the cleaning slot, you can set up a cleaning job for the robotic library drive.

See [“Cleaning a robotic library drive”](#) on page 567.

Note: Defined cleaning slots are not inventoried when an inventory job runs.

To define a cleaning slot

- 1 On the **Storage** tab, expand the robotic library, and then double-click **Slots**.
- 2 Double-click the slot that contains the cleaning tape.
- 3 In the **Cleaning slot** field, click the drop-down menu, and then click **Yes**.
- 4 Click **Apply**.

Editing robotic library properties

You can view robotic library properties.

See [“Robotic libraries in Backup Exec”](#) on page 478.

To view robotic library properties

- 1 On the **Storage** tab, double-click the robotic library for which you want to view properties.
- 2 In the left pane, click **Properties**.

3 Edit the following options as appropriate:

Name	Displays the name of the robotic library. You can edit this field.
Description	Displays the description of the robotic library. You can edit this field. By default, Backup Exec displays the device's inquiry string, which is the vendor name and product ID.
First slot number	Displays the starting slot for this robotic library. Backup Exec determines what the starting slot should be for this type of library. Some robotic libraries have slots that start at 0. Other libraries start at 1. You can change the starting slot if necessary.

4 Click **OK**.

Creating robotic library partitions

You can group one or more robotic library slots into partitions. Partitioning robotic library slots provides you with more control over which tape is used for backup jobs. When you create robotic library partitions, Backup Exec creates a storage device pool for each partition. Jobs that you send to a partition's storage device pool run on the media that is located in the partition's slots. For example, if you set up a partition that contains slots 1 and 2 and you want to run a weekly backup only on the media in these slots, you would submit the job to the storage device pool for the partition that contains slots 1 and 2. The storage device pools for robotic library partitions appear in the drop-down menu for the **Storage** field when you edit the backup job. All storage device pools for a robotic library partition have the same name and display the slot ranges for the partition in parentheses within the name.

Partitions can include any number of robotic library slots.

Depending on the robotic library configuration, the first slot may be numbered 1 or 0. If the robotic library uses a zero-based slot configuration, the **Configure Storage** wizard uses slot 0 as the first slot for partition 1 and adjusts the starting slot accordingly for all other partitions.

See ["Reassigning a slot base number for robotic libraries"](#) on page 488.

When the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the specified partition only. If more than one media is found that meets the requirements, Backup Exec then selects the media in the lowest-numbered slot; for example, media in slot 2 is selected before equivalent media in slot 4.

You can create a partitioning scheme that best suits your environment. For example, some administrators may create partitions based on users and groups, while others may create partitions according to operation types.

To create robotic library partitions

1 Do one of the following:

To view all of the robotic libraries that you can configure partitions for

Do the following in the order listed:

- On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- Select **Tape storage**, and then click **Next**.
- Select **Configure robotic library partitions**, and then click **Next**.
- Select the robotic library to configure partitions for, and then click **Next**.

To configure partitions for a specific robotic library

Do the following in the order listed:

- On the **Storage** tab, right-click the robotic library for which you want to create partitions.
- Click **Configure partitions**.

2 Specify the number of partitions to create, enter the number of slots for each partition, and then click **Next**.

3 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the partitions

Click **Finish**.

Adding or removing tape drives in a robotic library partition

You can add or remove tape drives in a robotic library partition.

See [“Creating robotic library partitions”](#) on page 486.

To add or remove tape drives in a robotic library partition

- 1 On the **Storage** tab, double-click the robotic library partition.
- 2 In the left pane, click **Properties**.
- 3 Do either of the following:

To add a tape drive to the robotic library partition	Do the following in the order listed:

- Click **Add**.
- In the devices list that appears, select the tape drive that you want to add, and then click **OK**.

To remove a tape drive from the robotic library partition	Select the tape drive that you want to remove, and then click Remove .

- 4 Click **Apply**.

Reassigning a slot base number for robotic libraries

Backup Exec automatically assigns slot base numbers for robotic libraries. If necessary, you can reassign how robotic library slots are displayed in Backup Exec. Slot base numbers in some robotic libraries start at 0, while slots in other robotic libraries start at 1. If the robotic library uses a zero-based slot configuration, you can reassign how the slots are displayed.

See [“Robotic libraries in Backup Exec”](#) on page 478.

To reassign a slot base number for robotic libraries

- 1 On the **Storage** tab, double-click the robotic library for which you want to reassign a slot base number.
- 2 In the **First slot number** field, click the drop-down menu to change the base number.
- 3 Click **Apply**.

Editing robotic library slot properties

You can edit the properties of a robotic library slot.

See [“Robotic libraries in Backup Exec”](#) on page 478.

To edit robotic library slot properties

- 1 On the **Storage** tab, double-click the robotic library.
- 2 In the left pane, click **Slots**.
- 3 Double-click the slot for which you want to view properties and then click **Properties**.
- 4 To view another slot's properties, click the drop-down menu at the top center of the window, and then click a slot number.
- 5 Edit any of the following options.

Cleaning slot

Indicates if this slot has been defined as a cleaning slot. If **Yes** is displayed, this slot has been defined as a cleaning slot.

Before you can submit a cleaning job, you must define a cleaning slot that contains the cleaning tape.

Make sure that the cleaning tape is located in the slot that you define as the cleaning slot. After defining the cleaning slot, you can set up a cleaning job for the robotic library drive.

Note: Defined cleaning slots are not inventoried when an inventory job runs.

See [“Cleaning a robotic library drive”](#) on page 567.

Preserve description

Keeps the media description when you select **Yes** in the drop-down menu. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

- 6 Click **Apply**.

Removing or reconfiguring robotic library partitions

You can remove or reconfigure robotic library partitions.

To remove robotic library partitions

- 1 On the **Storage** tab, right-click the robotic library that contains the partitions that you want to remove or reconfigure.
- 2 Click **Configure partitions**.
- 3 Click **Remove all robotic library partitions**, and then click **Next**.
- 4 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To remove the partitions

Click **Finish**.

To reconfigure robotic library partitions

- 1 On the **Storage** tab, right-click the robotic library that contains the partitions that you want to remove or reconfigure.
- 2 Click **Configure partitions**.
- 3 Click **Reconfigure robotic library partitions**, and then click **Next**.
- 4 Specify the number of partitions to create, enter the number of slots for each partition, and then click **Next**.
- 5 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the partitions

Click **Finish**.

See [“Reassigning a slot base number for robotic libraries”](#) on page 488.

Managing tapes

For tape media, you can perform the following actions:

- Protect data from being overwritten.
- Set up media rotation strategies.
- Track the location of media.
- Label media automatically.
- Read and track the media labels that have barcodes.
- Collect and report media statistics.

The Advanced Device and Media Management (ADAMM) feature in Backup Exec automatically selects the tape media for jobs. Backup Exec keeps track of all tape media that is loaded into the attached storage device. Backup Exec also keeps track of the media that is offline and the media that has been placed in media vaults.

For data that is kept on tapes, Backup Exec uses media sets to apply overwrite protection periods and append periods to manage the expiration of backup sets.

A media set consists of the following rules that apply to tape media:

- How long to protect the data on the media from overwrite. This is called the overwrite protection period.
- How long to append data to a media. This is called the append period.
- When and where to send media for vaulting.

Media that are associated with a media set are allocated media. Allocated media have current append and overwrite protection periods. Media that are associated with a media set but have expired overwrite protection periods are recyclable media.

For information about the best practices to manage tapes in Backup Exec, refer to *Backup Exec Best Practices*.

See [“Default media sets”](#) on page 491.

See [“Tape media rotation strategies”](#) on page 523.

See [“Creating media sets for tapes”](#) on page 499.

Default media sets

When you install Backup Exec, default system media sets and default user media sets are created automatically. When you add tapes to Backup Exec by importing media, Backup Exec associates the tape with one of the system media sets.

Note: You should not associate scratch media with a media set that you create. When a backup job runs, Backup Exec automatically moves the media from the Scratch Media set to the required media set as needed.

See [“Importing media to Backup Exec ”](#) on page 568.

You cannot modify the properties of system media sets. System media sets are described in the following table:

Table 13-2 Default system media sets

Name	Description
Backup Exec and Windows NT Backup Media	Displays all media that is imported from another installation of Backup Exec. See “Cataloging a storage device” on page 560.
Cleaning Media	Displays all cleaning media. See “Cleaning a robotic library drive” on page 567.
Foreign Media	Displays all media that is imported from a product other than Backup Exec. See “About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec” on page 253. See “Cataloging a storage device” on page 560.

Table 13-2 Default system media sets (*continued*)

Name	Description
Retired Media	Displays all media that you have taken out of service, usually because of an excessive number of errors. After you associate a media with the retired media set, Backup Exec does not select it for backup jobs. The media is still available for restore operations, if it has not been damaged. Retired Media protects media from being used (overwritten). If Backup Exec cannot recognize data on a tape, then it moves the tape to Retired Media. If you want to reuse the tape, erase or label the tape. These operations write a new header to the tape that Backup Exec can recognize. After the tape is erased or labeled, Backup Exec moves it to the Scratch Media set. You can delete the media that is in Retired Media to remove it from Backup Exec. You may want to delete media if you have a lot of off-site media that you do not want to recycle. You can also delete media if you throw away the media. See “Erasing tape or disk cartridge media” on page 518. See “Retiring damaged tape media” on page 517.
Scratch Media	Displays all media that can be overwritten. New, blank, and erased media are automatically associated with the Scratch Media set. See “Overwrite protection periods and append periods in media sets” on page 495.

Backup Exec creates the following default user media sets:

Table 13-3 Default user media sets

Name	Description
Keep Data for 4 Weeks	Displays all tape media that you associate with this media set. If you use the backup job defaults that are set when you install Backup Exec, the media set Keep Data for 4 Weeks is the default media set for all backup jobs that you send to tape storage. This media set protects data from being overwritten for four weeks and allows the media to be appended to for six days. You can edit and rename Keep Data for 4 Weeks after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data for 4 Weeks.
Keep Data Infinitely - Do Not Allow Overwrite	Displays all tape media that you associate with this media set. When you associate media with this media set, data is not overwritten unless you perform any of the following actions on the media: ■ Erase ■ Label ■ Format ■ Associate the media with the scratch media set You can append data to this media for an infinite period (until the media is full). You can edit and rename Keep Data Infinitely - Do Not Allow Overwrite after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data Infinitely - Do Not Allow Overwrite.

You can change the default media set for backup jobs by doing one of the following:

- Create new media sets that have the append and overwrite protection periods set to the time intervals that accommodate your data retention strategy. Then, specify the media set that is most appropriate when you create a backup job. For example, you can create a media set that keeps data for 60 days, and a media set that keeps data for 90 days.
- Select the other default media set **Keep Data Infinitely - Do Not Allow Overwrite** when you create a backup job. The risk that is associated with the media set **Keep Data Infinitely - Do Not Allow Overwrite** is that you can

use all of your scratch media. You must continually add new tape or disk cartridge media to Backup Exec.

Note: It is recommended that if you need to keep data longer than four weeks, you should duplicate it. You can duplicate the backup data from the original storage device to tape, which you can then send for long-term or off-site storage.

To view all media sets

◆ On the **Storage** tab, double-click **All Media Sets**.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

See [“Creating media sets for tapes”](#) on page 499.

See [“Creating media vault rules to move tape media to and from media vaults”](#) on page 515.

See [“Duplicating backup sets or a job history manually”](#) on page 222.

See [“Associating tape media with a media set”](#) on page 520.

Overwrite protection periods and append periods in media sets

Each tape media is associated with a media set, which is a set of rules that manage media.

These rules include overwrite protection and append periods.

Table 13-4 Overwrite protection and append periods

Rule	Description
Append period	The amount of time that data can be appended to tape media. It is measured from the time the media was first allocated. It can be specified in hours, days, weeks, or years.

Table 13-4 Overwrite protection and append periods (*continued*)


Rule	Description
Overwrite protection period	The amount of time that tape media is protected from being overwritten. The period is measured from the time of the last write to the media, that is, at the end of the last append or overwrite job. It can be specified in hours, days, weeks, or years. When the overwrite protection period is over, the media becomes recyclable and can be overwritten. The overwrite protection period begins when the backup job is completed. If there is an append period, the overwrite protection period begins again each time an append job completes. Because the overwrite protection period does not begin until the job completes, the amount of time that is required to complete the job affects when the media can be overwritten. You may shorten the overwrite protection period to take into account the amount of time a job may run. For example, you set the overwrite protection period for seven days. You also set the append period for four days to ensure that data is not overwritten for at least seven days. The data can be appended to the media for the next four days. The last data that is appended to this media is retained for at least seven days. Note: Any media can be overwritten if the overwrite protection level is set to None. See “Media overwrite protection levels for tape media” on page 505.

Your tape media rotation strategy must balance between the need to save data as long as possible, and the fact that tape media are not in infinite supply. The media set rules allow Backup Exec to identify which tape media can be written to and which tape media are overwrite-protected. You should consider the use of disk storage for backup data.

See [“Configuring disk storage”](#) on page 334.

The following graphic shows the relationship between the append period and the overwrite protection period.


Figure 13-1 Append periods and overwrite protection periods



The append and overwrite protection periods that you specify apply to all the data on the media.

Each time data is written to a media, the time remaining in the overwrite protection period is reset, and the countdown is restarted.

Figure 13-2 How overwrite protection periods are reset



The amount of time that is required to complete the job affects when the media can be overwritten.

For example, suppose that you create a media set named *Weekly* with an overwrite protection period of seven days. You also specify an append period of zero days, and you schedule a full backup job to run each Friday at 20:00. When it is time for the full backup to run at 20:00 the following Friday, the job cannot run. The first backup job that ran the previous Friday did not complete until 21:10. The overwrite protection period for the *Weekly* media set still has 70 minutes remaining.

To prevent this situation, you can shorten the overwrite protection period to account for the amount of time a job may run. For this example, the scheduled job recurring at 20:00 can run if the overwrite protection period is set to six days instead of seven days.

Creating media sets for tapes

A media set consists of the rules that specify the following:

- Append periods
- Overwrite protection periods
- Media vaults
- Amount of time to move media to and from the media vault

The media set rules apply to all of the media that you associate with the media set.

Note: You must have already created a media vault before you are prompted to add media vault rules in a media set.

See [“Default media sets”](#) on page 491.

To create a media set for tapes

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Under **User Media Sets**, right-click a media set, and then click **Create media set**.
- 4 Enter a name and a description for the media set, and then click **Next**.
- 5 Specify how long you want to keep data in this media set from being overwritten, and then click **Next**.
- 6 Specify how long you want to allow data to be appended to media in this media set, and then click **Next**.
- 7 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To create the media set

Click **Finish**.

Editing media set properties

You can edit the following properties for media sets:

- Name of a media set.
- Overwrite protection and append periods for a media set.
- Media vault and the vaulting periods associated with a media set.

See [“Default media sets”](#) on page 491.

To edit media set properties

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to edit properties.
- 4 In the left pane, click **Properties**.
- 5 Change any of the following information:

Name	Displays the name of the media set.
Description	Displays the description of the media set.

Overwrite protection period

Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.

Regardless of the overwrite protection period that is set, media can be overwritten if you perform the following operations on it:

- Erase
- Label
- Associate it with the **Scratch Media Set**
- Set the **Media Overwrite Protection Level** to **None**
- Format

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The default period is **Infinite - Don't Allow Overwrite**, which protects the media from being overwritten for 1,000 years.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

Append period

Displays the length of time in hours, days, or weeks that data can be added to media. Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The append period starts when the first backup job is written to this media.

The default period is **Infinite - Allow Append**, which allows data to be appended until the media capacity is reached.

Media vault to use with this media set	Displays the media vault that stores the media that is associated with this media set. See “Default media vaults” on page 513.
Move media to this vault after	Displays the time period after which this media is reported as ready to be moved to this vault.
Return media from this vault after	Displays the time period after which this media is reported as ready to be returned from this vault.

- 6 Click **Apply**.

Changing the overwrite protection period or the append period for a media set

You can change the length of time to retain data on media that are associated with a media set. You can also change the length of time to add data to media that are associated with a media set.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

To change the overwrite protection period or the append period for a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the overwrite protection period or the append period.
- 4 In the left pane, click **Properties**.
- 5 Change the following information as appropriate:

Overwrite protection period

Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.

Regardless of the overwrite protection period that is set, media can be overwritten if you perform the following operations on it:

- Erase
- Label
- Associate it with the **Scratch Media Set**
- Set the **Media Overwrite Protection Level** to **None**
- Format

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The default period is **Infinite - Don't Allow Overwrite**, which protects the media from being overwritten for 1,000 years.

Append period

Displays the length of time in hours, days, or weeks that data can be added to media.

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The append period starts when the first backup job is written to this media.

The default period is **Infinite - Allow Append**, which allows data to be appended until the media capacity is reached.

Changing the name and description of a media set

You can change the name and description of a media set.

To change the name and description of a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the name or description.
- 4 In the left pane, click **Properties**.
- 5 Change the name and/or the description of the media set, and then click **Apply**.

Changing the media vault or vaulting periods for a media set

You can change the media vault that stores this media set, and you can change the time periods for when you want to move media to a vault or return media from a vault.

See [“Default media vaults”](#) on page 513.

To change the media vault or vaulting periods for a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the media vault or vaulting periods.
- 4 In the left pane, click **Properties**.
- 5 Change the following information as appropriate:

Media vault to use with this media set	Displays the media vault that stores the media that is associated with this media set.
Move media to this vault after	Displays the time period after which this media is reported as ready to be moved to this vault.
Return media from this vault after	Displays the time period after which this media is reported as ready to be returned from this vault.

- 6 Click **Apply**.

Deleting a media set

If you delete a media set that has scheduled jobs associated with it, you are prompted to associate the jobs to another media set.

Caution: Ensure that the media set that you associate the jobs with has the appropriate overwrite protection and append periods.

See [“Default media sets”](#) on page 491.

To delete a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Right-click the media set that you want to delete, and then click **Delete**.
- 4 When you are prompted to delete the media set, click **OK**.

Renaming a media set

When you rename a media set, any jobs that use that media set display the new media set name.

See [“Default media sets”](#) on page 491.

To rename a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Right-click the media set that you want to rename, and then click **Details**.
- 4 In the left pane, click **Properties**.
- 5 In the **Name** field, type the new name that you want to assign to this media set, and then click **Apply**.

Media overwrite protection levels for tape media

The media overwrite protection level is a global setting that supersedes the media set's overwrite protection period. Although the terms are similar, the media overwrite protection level and the media overwrite protection period are different. The media overwrite protection period is a time interval that changes from one media set to another. The media overwrite protection level specifies whether to overwrite scratch, imported, or allocated media, regardless of the media's overwrite protection period.

Use the media overwrite protection level to specify the type of media that you want to be available for overwrite backup jobs.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

Overwriting allocated or imported tape media

Media that are associated with a media set are called allocated media. Media that are imported from another installation of Backup Exec, or from another product are called imported media. Backup Exec protects allocated and imported media from being overwritten when full or partial overwrite protection is used. However, you can let Backup Exec overwrite allocated and imported media before the data overwrite protection period expires, and without setting the media overwrite protection level to None.

The following methods are available:

- Associate the media with the **Scratch Media Set**. The media is overwritten when it is selected for an overwrite job.
- Erase the media. Erased media is automatically recognized as scratch media and is overwritten immediately.
- Label the media. The **Label Media** operation immediately writes a new media label on the media, which destroys any data that is contained on the media.
- Format the media. Formatting destroys any data that is contained on the media.
- Change the overwrite protection period for the media set so that it is expired.

See [“Managing tapes”](#) on page 491.

See [“Deleting tape media”](#) on page 518.

See [“Editing global settings for storage”](#) on page 548.

How Backup Exec searches for overwritable media in tape drives

Media overwrite options set the order in which Backup Exec searches for overwritable media in tape drives. When Backup Exec searches for overwritable media for a backup job, it searches for either scratch media or media that has an expired overwrite protection period.

You are prompted to select one of the following options that you want Backup Exec to use first:

- Overwrite scratch media before overwriting recyclable media that is contained in the destination media set.

If you choose to overwrite scratch media before recyclable media, more media may be required for the same number of jobs. However, the recyclable media may be preserved longer for possible recovery.

- Overwrite recyclable media that is contained in the destination media set before overwriting scratch media.

If you choose to overwrite recyclable media before scratch media, the same media is re-used more frequently than if you overwrite scratch media before recyclable media.

In a storage device pool for tape drives, Backup Exec selects the oldest recyclable media in the storage device pool to use first.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the targeted partition only.

Caution: It is recommended that you physically write-protect media containing critical data. Use the write-protect tab on the media cartridge to protect against unintentional move or erase operations, or expired overwrite protection periods.

The following table describes the order in which Backup Exec searches for media to use for an overwrite job.

Table 13-5 How Backup Exec searches for overwritable media in tape drives

Overwrite protection level and overwrite option:	Media is overwritten in tape drives in this order:	
Full + Overwrite scratch media first Note: This combination provides the most protection against overwriting media.	1	Scratch media
	2	Recyclable media in the destination media set
	3	Recyclable media in any media set
Full + Overwrite recyclable media first	1	Recyclable media in the destination media set
	2	Scratch media
	3	Recyclable media in any media set
Partial + Overwrite scratch media first	1	Scratch media
	2	Recyclable media in the destination media set
	3	Recyclable media in any media set
	4	Media that is imported from another installation of Backup Exec, or from another product

Table 13-5 How Backup Exec searches for overwritable media in tape drives
(continued)

Overwrite protection level and overwrite option:	Media is overwritten in tape drives in this order:	
Partial + Overwrite recyclable media first	1	Recyclable media in the destination media set
	2	Scratch media
	3	Recyclable media in any media set
	4	Media that is imported from another installation of Backup Exec, or from another product
None - No overwrite protection + overwrite scratch media first Warning: This options is not recommended because it does not protect data from being overwritten.	1	Scratch media
	2	Recyclable media in the destination media set
	3	Recyclable media in any media set
	4	Media that is imported from another installation of Backup Exec, or from another product
	5	Allocated media in any media set
None - No overwrite protection + overwrite recyclable media first Warning: This options is not recommended because it does not protect data from being overwritten.	1	Recyclable media in the destination media set
	2	Scratch media
	3	Recyclable media in any media set
	4	Media that is imported from another installation of Backup Exec, or from another product
	5	Allocated media in any media set

In addition to setting overwrite protection levels, you must set overwrite options, which set the order in which Backup Exec searches for overwritable media.

The most obvious candidates for backup jobs that require overwritable media are scratch media and recyclable media. Recyclable media have expired overwrite protection periods. Backup Exec searches for these types of media first when a backup requires tape media to overwrite. The search pattern is different according to whether you have selected Full, Partial, or None. The media indicate that a type of media set is examined for availability.

See [“Editing global settings for storage”](#) on page 548.

See [“Default media vaults”](#) on page 513.

Viewing tapes that are used by a job

To see which tapes are used for a job, review the Device and Media Information section of the job log.

See [“Managing tapes”](#) on page 491.

To view the tapes that are used by a job

- 1 Do any of the following:

To view the job log from the **Job Monitor** tab Select the **Job Monitor** tab.

To view the job log from the **Backup and Restore** tab or the **Storage** tab Do the following in the order listed:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and the select **View Job Log**.
- 3 Review the information in the section **Device and Media Information**.

Labeling tape media

Media labels identify the tapes that you use in Backup Exec. When a new, blank, or unlabeled tape is used during a backup operation, Backup Exec automatically labels the tape media. This label consists of a prefix that identifies the cartridge type, and an incrementing number. For example, if the media is a 4mm tape, then the prefix is 4M, followed by 000001. The next media label generated for an unlabeled 4mm tape would be 4M000002, and so on.

Another type of media label used by Backup Exec is the media ID, which is a unique label assigned by Backup Exec to the individual tape media used in Backup Exec. The media ID is used internally by Backup Exec to keep statistics on each media. Because the media label or barcode label for tape media can be changed, Backup Exec must use the media ID to preserve continuity in record keeping for each individual tape media. You cannot change or erase the media ID. The media ID has no effect on the media label, or on your ability to rename, label, or erase tape media.

At times, you may need to use the media ID to distinguish the tape media that have duplicate media labels. Duplicate labels can be automatically generated in instances when Backup Exec is reinstalled or media from another Backup Exec installation is used. You can view the media ID in a media's property page.

Write the media label on an external label that is fixed to the outside of the physical tape cartridge. Whenever you change the media label, you should also change the external label to match.

The following methods are available in Backup Exec to change a tape media label:

- Write a new media label on the tape media. The Label operation destroys any data on the media.
- Rename the media. Renaming the tape media changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.
- Edit the label. Editing the label changes the name of the tape media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.

Note: Media that use barcode labels cannot be renamed. When you try to label the media that use barcode labels, the job logs report successfully completed jobs. However, the media label names do not change.

To label media

- 1 On the **Storage** tab, do either of the following:
 - Right-click the drive that contains the tape that you want to label.
 - Double-click **Slots**, and then right-click the slot that contains the tape that you want to label.
- 2 Click **Label**.

The following warning appears:

This operation is performed on the current media in the drive or slot. If the media has changed since the last inventory ran, the media label in the next dialog may not match the media in the selected device.
- 3 Click **OK**.
- 4 Type the name that you want to use as the media label for this media.
- 5 To erase all data on the media and re-label the media, click **OK**.
- 6 Write this same media label on an external label that is fixed to the outside of the physical media.
- 7 (Optional) View the job history for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

See [“Renaming a tape media label ”](#) on page 512.

See [“How barcode labels become media labels for tapes”](#) on page 511.

About labeling imported tape media

Tape media that is imported from another installation of Backup Exec, or from another product is called imported media. Backup Exec does not automatically relabel imported media.

Backup Exec reads the imported tape media’s existing label and displays the label in either the **Backup Exec and Windows NT Backup Media** media set or in the **Foreign Media** media set. If the media overwrite protection level is set to Partial or None, the imported media may be selected for a job and be overwritten. The imported media is automatically labeled when it is overwritten during a job. If you want to label a specific imported media while maintaining full media overwrite protection for other imported media, erase the specific media and then label it.

The original media label of the imported tape media is displayed in the media’s properties. You can edit the media description in the media’s property page to make it a more descriptive label.

See [“Erasing tape or disk cartridge media”](#) on page 518.

How barcode labels become media labels for tapes

If there is a barcode label on the physical tape cartridge, and if the robotic library has a barcode reader, the barcode label automatically becomes the media label.

For example, robotic library 1 has barcode support. During a backup operation, Backup Exec requests a new media or an overwritable media for the operation. A new media with the barcode label 'ABCD' is inserted in the robotic library magazine and the barcode reader scans the barcode ID. Backup Exec selects this media for the operation and detects that a barcode label has been assigned to the media. Backup Exec automatically uses the barcode label and continues the operation.

If a barcode label is the media label, then you cannot change the media label in Backup Exec. To change a barcode label and use a media label, you must remove the physical barcode label from the media cartridge. Or, you can use the media in a device that does not have a barcode reader.

When you change magazines or insert new media in a magazine, you can use the **Scan** operation to quickly update slot information.

See [“Labeling tape media”](#) on page 509.

See [“Scanning a storage device”](#) on page 561.

Renaming a tape media label

You can rename a tape media's label and description. The new label is not written to the tape media until an overwrite operation occurs. All of the data that is on the media is preserved until the next overwrite job. However, the new media label is stored in the database and is displayed for that media. To write a new media label to the media immediately, use the **Label** operation. The media's contents are erased.

If you rename a tape media, and then use it in another installation of Backup Exec, that media is imported to the **Backup Exec and Windows NT Media** media set. The media's original media label is displayed. The renamed label is not transferred to other installations of Backup Exec.

Note: If a barcode label is the media label, then you cannot change the media label in Backup Exec.

See [“Labeling tape media”](#) on page 509.

To rename a tape media label

- 1 On the **Storage** tab, right-click the tape drive or slot that contains the media that you want to relabel, and then click **Details**.
- 2 In the left pane, click **Media Properties**.
- 3 In the **Media label** field, enter a new label name.
- 4 To change the description, enter a new description in the **Media description** field.
- 5 Click **Apply**.

How WORM media is used in Backup Exec

Write once, read many (WORM) data storage is used to keep the data that has a long retention period. Data can be written to WORM media one time only. After it is written to, the media can be appended to, but it cannot be overwritten, erased, or reformatted.

When WORM media is used in a media set, the overwrite protection period is not applied to it, but the append period is applied.

New WORM media is WORM media that has not been written to. When new WORM media is introduced into Backup Exec, it is placed in the **Scratch Media**

set. After the WORM media has been written to one time, you cannot move it to the Scratch media set. You can move WORM media to the **Retired Media** set to delete it from Backup Exec, but you cannot erase it or reformat it.

You can format a tape as a WORM tape if the tape drive supports the operation.

See [“Formatting a tape as a WORM tape”](#) on page 565.

When you select the option to use WORM media, Backup Exec verifies that the destination device is or contains a WORM-compatible drive. Backup Exec also verifies that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent and the backup job may fail.

See [“Configuring storage options for backup jobs”](#) on page 652.

Default media vaults

A media vault is a logical representation of the actual physical location of specified tape media. You can create media vaults to keep track of where media is physically stored, such as a scratch bin, or an off-site location. Backup Exec creates default media vaults to let you view all media that are online, offline, or in a media vault.

You must run the **Configure Storage** wizard to update the location of tape media in media vaults. From the **Configure Storage** wizard, you can print reports that detail which tape media are ready to move to and return from the vault. You can also update the location of the media if you choose to move them. However, you must physically collect the tape media, and move the media to and from the vault. The location of the tape media is updated in the Backup Exec Database, but the media is not ejected or exported. If Backup Exec detects that the media is in a robotic library, you are prompted to export media. If you choose to export the media, an export media job runs. If your environment includes remote sites, you should create separate media sets for each remote site. Then, the reports contain details on which media are ready to be moved for a specific site.

Table 13-6 Default media vaults

Default media vault	Description
Online Tape Media	Displays the media that are available in tape drives or robotic libraries. You cannot add or move media to the online media vault. Backup Exec does that automatically. If you move media from the online media vault to another media vault, the media's overwrite protection period and append period remain in effect.

Table 13-6 Default media vaults (continued)

Default media vault	Description
Offline Tape Media	Displays the media that are on-site but are not in tape drives or robotic libraries, and are not in media vaults. Media appear in the offline media vault if you use Backup Exec to remove media from a tape drive or robotic library. You can add media to the offline media vault from another media vault. An inventory operation or a catalog operation moves the offline media back to the online media vault. You cannot delete or rename the offline media vault.
Vaulted Tape Media	Displays the media that are not in tape drives or robotic libraries, and have been moved to a media vault. Vaulted Tape Media displays in All Media Vaults details only after you create a media vault.
All Media Vaults	Displays the media that are in media vaults that you create. All Media Vaults displays on the Storage tab only after you create a media vault. You can associate media vaults with media sets that you create. You specify when to move the media from a media set to the media vault. You also specify when the media is to return to the media set from the media vault. See “Creating media vault rules to move tape media to and from media vaults” on page 515.

See [“Changing the name or description of a media vault”](#) on page 514.

Changing the name or description of a media vault

You can edit the name and description of a media vault.

See [“Default media vaults”](#) on page 513.

To change media vaults and media vault rules, edit the properties of the media set that is associated with the media vault.

See [“Editing media set properties”](#) on page 500.

Changing the name or description of a media vault

- 1 On the **Storage** tab, expand **Tape Cartridge Media Sets and Vaults**, and then expand **All Media Vaults**.
- 2 Right-click the media vault for which you want to edit properties, and then click **Details**.
- 3 In the left pane, click **Media Vault Properties**.
- 4 Change the name or description of the media vault.
- 5 Click **Apply**.

Media vault properties

Properties for media vaults include the name and a description of the media vault.

See [“Changing the name or description of a media vault”](#) on page 514.

Table 13-7 Properties for media vaults

Item	Description
Name	Displays the name of the media vault.
Description	Displays a description of the media vault.

Creating media vault rules to move tape media to and from media vaults

Create media vault rules to do the following:

- Associate a media vault with the media set to which you want to send media.
- Specify the amount of time to wait between when the media is allocated and when it is sent to the vault.
- Specify the amount of time to wait between returning the media from the vault and when it was last written to.

See [“Default media vaults”](#) on page 513.

Backup Exec does not update the vault automatically. You must use the **Configure Storage** wizard to update the location of the tape media. You can also print or view the reports that contain details on which media are ready to be moved to and from the vault.

See [“Updating the tape media location in media vaults”](#) on page 516.

To create media vault rules to move tape media to and from media vaults

- 1 On the **Storage** tab, expand **All Media Sets**.
- 2 Right-click **Keep Data for 4 Weeks, Keep Data Infinitely - Do Not Allow Overwrite**, or a media set that you created, and then click **Details**.
- 3 In the left pane, click **Properties**.
- 4 Select the media vault that you want to use with the media set.
- 5 Specify when to move the media to the vault and when to return the media to the media set.

Updating the tape media location in media vaults

You can update the location of tape media that are in vaults. You can also print the reports that detail which media are ready to move to and return from the vault. However, you must physically collect the media, and move the media to and from the vault.

See [“Default media vaults”](#) on page 513.

To update the tape media location in media vaults

- 1 On the **Storage** tab, expand **Tape Cartridge Media Sets and Vaults**, and then double-click **All Media Vaults**.
- 2 Right-click the media vault for which you want to update the media location, and then click **Update vault using wizard**.
- 3 Follow the on-screen prompts.

Deleting a media vault

You can only delete an empty media vault. If any tape media in the vault, you must move it before you can delete the vault. You cannot delete the online media vaults or the offline media vaults.

See [“Default media vaults”](#) on page 513.

To delete a media vault

- 1 On the **Storage** tab, expand **All Media Vaults**.
- 2 Right-click the media vault that you want to delete, and then click **Delete**.
- 3 Click **Yes** when you are prompted to delete the media vault.

Moving tape media to a vault

You can use a barcode scanner to enter the media labels of tape media that you want to move to a vault. You can also type a media label into the dialog box.

See [“Default media vaults”](#) on page 513.

To move tape media to a vault

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**, and then double-click **All Media Vaults**.
- 2 Right-click the media vault that you want to move media to, and then click **Move media to vault**.
- 3 Follow the on-screen prompts.

Retiring damaged tape media

You can retire damaged tape media so that Backup Exec does not use it for backup jobs. You should associate tape media that meets or exceeds the discard thresholds that are determined by the media manufacturer with the **Retired Media** media set. Backup Exec tracks the soft errors that are generated by the storage device firmware. Media that exceed acceptable levels of these errors are reported as potential candidates to be discarded.

To decide which tape media to retire, run a Media Errors report for the total number of errors for media, or view the properties for a specific media.

Associate any media with an unacceptable level of errors to **Retired Media** so that you are protected against using defective media before critical backup operations begin. After you associate tape media with the **Retired Media** set, it is not used by Backup Exec for future backup jobs. The media is still available to be restored from if it is not damaged.

To retire damaged tape media

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Right-click **Online Tape Media**, and then click **Details**.
- 3 Right-click the media that you want to retire, and then click **Retire**.
- 4 Click **Yes** when you are prompted to retire the media.

See [“Deleting tape media”](#) on page 518.

See [“Media Errors report”](#) on page 825.

Deleting tape media

When you delete tape media from Backup Exec, all records of the media are removed from the Backup Exec database. These records include catalog information, media statistics, and other information that is associated with the tape media. You can only delete media when it belongs to the **Retired Media** set.

You may want to delete tape media when the following occurs:

- You have a lot of off-site media that you do not want to recycle.
- You throw away damaged or old media.

If you import deleted media back into Backup Exec, it is added to either the **Backup Exec and Windows NT Media** media set or the **Foreign Media** media set. Before you can restore from the media, you must catalog it.

To delete tape media

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Expand **All Media Sets**, right-click **Retired Media**, and then click **Details**.
- 3 Right-click the media that you want to delete, and then click **Delete**.
- 4 Click **Yes** when you are prompted to delete the media.

See [“Retiring damaged tape media”](#) on page 517.

Erasing tape or disk cartridge media

You can erase tape media or disk cartridge media immediately, or you can schedule the erase operation.

Warning: The erase operation is run on whatever media is in the drive or slot at the time when the operation runs. If the media has been changed since the last inventory operation was run, the media label that appears on the Backup Exec Administration Console may not match the media in the selected drive or slot. If media is moved unexpectedly, data loss can occur. Check any scheduled erase jobs carefully.

Table 13-8 Erase operations

Erase operation	Description
Erase media now	Writes an indicator at the beginning of the media that makes the data that is contained on the media inaccessible. For most uses, the Erase media now operation is sufficient. This is the only erase operation available for disk cartridge media.
Long erase media now	Instructs the drive to physically erase the entire media. If you have sensitive information on the media and you want to dispose of it, use the Long erase media now operation. A long erase operation on a media will take several hours to complete depending on the drive and the media capacity. Some devices do not support a long erase operation.
Schedule	Let you schedule either an Erase or Long erase operation, and lets you choose notification options.

The erase operation does not change the media label. To change a tape media label, either run a Label operation or rename the media before you run an Erase operation.

You cannot cancel an Erase operation after it has started. You can cancel an Erase operation that is scheduled or queued.

When you schedule an erase operation, you can configure the time and frequency that you want to run the job.

To erase tape or disk cartridge media now

- 1

On the **Storage** tab, right-click the drive or the robotic library slot that contains the media that you want to erase.
- 2

Click **Erase media now**, and then do one of the following:

To run an erase operation immediately

Click **Erase media now**.

To run a long erase operation immediately

Click **Long erase media now**.
- 3

Click **Yes** when you are prompted to erase the media.

To schedule an erase operation for tape or disk cartridge media

- 1 On the **Storage** tab, right-click the drive or the robotic library slot that contains the media that you want to erase.
- 2 Click **Erase media now**, and then do one of the following:

To schedule an erase operation	Click Schedule erase .
To schedule a long erase operation	Click Schedule long erase .
- 3 Click **Yes** when you are prompted to erase the media.
- 4 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.
See [“Scheduling a storage operation job”](#) on page 541.
- 6 Click **OK**.

About cataloging tape or disk cartridge media that contains encrypted backup sets

When you catalog tape media or disk cartridge media that contains encrypted backup sets, Backup Exec attempts to find valid encryption keys for the sets in the Backup Exec database. If Backup Exec does not find a valid key, it issues an alert that instructs you to create one. After you create a valid key, you can respond to the alert to retry cataloging the encrypted set. Alternatively, you can skip the encrypted set and continue to catalog the rest of the media, or cancel the catalog job.

See [“Encryption key management”](#) on page 733.

Associating tape media with a media set

When you create a backup job to tape media, the default media set that Backup Exec selects for you is called **Keep Data for 4 Weeks**. You can select other media sets when you create the backup job, or you can associate the tape media with another media set later.

When you associate tape media with a media set, the tape uses the following properties of that media set:

- Append periods
- Overwrite protection periods
- Media vaults
- Amount of time to move media to and from the media vault.

Note: You should not associate scratch or imported media with a media set. Backup Exec automatically associates scratch or imported media with a media set when a backup job requires it.

See [“Creating media sets for tapes”](#) on page 499.

To associate tape media with a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Tape Media** to display a list of media.
- 3 Right-click the tape media that you want to associate with a media set, and then click **Associate with media set**.
- 4 Select a media set from the drop-down list, and then click **OK**.

See [“Default media sets”](#) on page 491.

Editing tape media properties

You can view tape media properties and edit some properties.

See [“Managing tapes”](#) on page 491.

To edit tape media properties

- 1 On the **Storage** tab, double-click the drive that contains the media.
- 2 In the left pane, click **Media Properties**.

3 Change any of the following options:

Media label

Displays the media label that Backup Exec assigns automatically, or that the administrator assigned, or that is a pre-assigned barcode label.

You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. When you edit a media label, try to make it a concise identifier that remains constant even when the media is reused. You should write this media label on a label that is fixed to the outside of the physical media.

Duplicate labels can be automatically generated. For example, reinstalling Backup Exec or bringing media from another Backup Exec installation can cause duplication in labels. Duplicate labels are allowed, but not recommended.

If a barcode is available, and a barcode-equipped device is used, then the media label automatically defaults to that barcode.

Media description

Displays the original media label if the media is imported media.

You can edit the media description, which is limited to 128 characters, to make it a more descriptive label.

Preserve description

Keeps the media description when you select **Yes** in the drop-down box. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

4 Click **Apply**.

Tape media rotation strategies

Many tape media rotation strategies exist that you can use to back up your data.

The most commonly used tape media rotation strategies include the following:

- Son, which uses the same tape each day to run a full backup.
- Father/Son, which uses multiple tapes, and includes a combination of weekly full and daily differential or incremental backups for a two-week schedule. This strategy provides backups for off-site storage .
- Grandfather, which uses multiple tapes, and includes a combination of weekly and monthly full and daily differential or incremental backups. This strategy also provides backups for off-site storage.

Son media rotation strategy


The Son media rotation strategy requires the following:

Table 13-9 Son media rotation strategy

Item	Description
Number of media required	1 (minimum)
Overwrite protection period	Last backup

The Son strategy involves performing a full backup every day.

Figure 13-3 Son backup strategy



Although the Son strategy is easy to administer, backing up with a single media is not an effective method of backup. Magnetic media eventually wears out after many uses and the data you can restore only spans back to your last backup.

Father/Son media rotation strategy

The Father/Son media rotation strategy requires the following:

Table 13-10 Father/Son media rotation strategy

Item	Description
Number of media required	6 (minimum)
Overwrite protection period	Two weeks


The Father/Son media rotation strategy uses a combination of full and differential or incremental backups for a two-week schedule.

In the Father/Son scenario, four media are used Monday through Thursday for differential or incremental backups. The other two media containing full backups are rotated out and stored off-site every Friday.

The Father/Son strategy is easy to administer and lets you keep data longer than the Son strategy. The Father/Son strategy is not suitable for the stringent data protection needs of most network environments.

When this backup strategy is first implemented, you must start with a full backup.

Figure 13-4 Father/Son backup strategy



Grandfather media rotation strategy

The Grandfather media rotation strategy requires the following:

Table 13-11 Grandfather media rotation strategy


Item	Description
Number of media required	19 (minimum)
Overwrite protection period	One year

The Grandfather method is one of the most common media rotation strategies. The Grandfather method is easy to administer and comprehensive enough to allow easy location of files when they need to be restored.

In the Grandfather scenario, four tapes are used Monday through Thursday for incremental or differential backups; another three tapes are used every Friday for full backups.

The remaining 12 media are used for monthly full backups and are kept off-site.

Figure 13-5 Grandfather backup strategy



The Grandfather strategy is recommended because it offers a good media number to storage life ratio (19 media/1 year). You can easily incorporate more media. For example, you can perform a full backup on the last Saturday of the month to keep permanently.

Storage device pools

This chapter includes the following topics:

- [Creating storage device pools](#)
- [Specifying a default selection method for a device in a storage device pool](#)
- [Viewing jobs, job histories, and active alerts for a storage device pool](#)
- [Adding or removing devices in a storage device pool](#)

Creating storage device pools

A storage device pool is a group of similar types of storage devices that enables load-balancing of Backup Exec jobs. The workload is shared across the storage device pool. You can send backup jobs to specific storage devices or to a storage device pool. By default, if the specific storage device is busy, the job must wait until the storage device becomes available.

When you send a job to specific storage devices, Backup Exec cannot automatically route the job to the next available storage device. When you submit a backup job to a storage device pool, by default the job is sent to the first available storage device in that pool. As other jobs are created and started, they can run concurrently on other storage devices in the storage device pool. By dynamically allocating storage devices as jobs are submitted, Backup Exec processes jobs quickly and efficiently. Storage device pools provide fault tolerance if you configure error-handling rules to resubmit a job that fails because of a storage device error.

You can change the default selection method for a device in a storage device pool so that jobs are sent first to one of the following devices in the storage device pool:

- A storage device with the least amount of free space. This setting is beneficial for manual offsite rotation strategies since the job runs to the same device until the device is full.
- A storage device with the most amount of free space. This setting is beneficial for a backup that has Granular Recovery Technology (GRT) enabled, and is sent to a disk storage device pool. GRT jobs that are sent to disk storage devices cannot span.

Note: You must use the Backup Exec Management Command Line Interface to change the default to specify whether jobs are sent to a device that has the most or least free space.

See [“Specifying a default selection method for a device in a storage device pool”](#) on page 530.

Backup Exec creates and maintains system-defined storage device pools for disk storage, tape drives, disk cartridges, and virtual disks. Any storage devices that you configure or that you attach are automatically added to the appropriate system-defined storage device pool. You can select these storage device pools when you select the storage for a backup job. You cannot edit a system-defined storage device pool. Before you can view the default storage device pools on the **Storage** tab, under **All Storage Pools**, you must create a storage device pool. If you have the Central Admin Server feature installed, you can create managed Backup Exec server pools.

Table 14-1 System-defined storage device pools maintained by Backup Exec

System-defined storage device pools that are maintained by Backup Exec	Description
Any disk storage	Contains the fixed-disk storage.
Any tape drive	Contains the tape drives. Backup Exec creates this pool when it detects an attached tape drive or robotic library. In a tape drive storage pool, Backup Exec uses the oldest recyclable media first.
Any disk cartridge	Contains the disk cartridges that you have configured to use as storage. Backup Exec creates this pool the first time that you configure disk cartridge storage.

For storage device pools that you create, you must edit the properties of the pool and check the check box next to the device before jobs can use that device in the pool.

Use the **Configure Storage** wizard to create storage device pools.

To create a storage device pool

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server feature is not installed Select **Storage pools**, and then click **Next**.

If the Central Admin Server feature is installed

Do the following in the order listed:

- Select the Backup Exec server that you want to configure storage for, and then click **Next**.
- Select **Storage pools**, and then click **Next**.

- 3 Select **Storage device pool**, and then click **Next**.
- 4 Enter a name and description for the pool, and then click **Next**.
- 5 Click the drop-down arrow, select the type of storage device pool that you want to configure, and then click **Next**.
- 6 Select all of the storage devices that you want to include in the pool, and then click **Next**.
- 7 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the storage device pool

Click **Finish**.

See [“Adding or removing devices in a storage device pool ”](#) on page 531.

See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.

See [“Adding or removing devices in a storage device pool ”](#) on page 531.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

Specifying a default selection method for a device in a storage device pool

The default selection method for a storage device pool is the first available storage device in a pool. You can use the Backup Exec Management Command Line Interface to specify another default selection method for a disk-based storage device pool. You can change the default setting of an existing storage device pool, or a storage device pool that Backup Exec creates.

You can change the default so that jobs are first sent to one of the following devices:

- A storage device with the least amount of free space. This setting is beneficial for manual off-site rotation strategies since the job runs to the same device until the device is full.
- A storage device with the most amount of free space. This setting is beneficial for a backup job for which Granular Recovery Technology (GRT) is enabled, and that is sent to a disk storage device pool. A GRT job cannot span when it is sent to a disk storage device.

The Backup Exec Management Command Line Interface commands are as follows:

- `New-BESStorageDevicePool` sets the device selection method for a new disk-based storage device pool when you create it.
- `Set-BESStorageDevicePool` sets the device selection method for an existing disk-based storage device pool.

For information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:<Backup Exec install path>\Backup Exec

See [“Creating storage device pools”](#) on page 527.

Viewing jobs, job histories, and active alerts for a storage device pool

You can view the jobs that are sent to a storage device pool, and the job histories as well as any active alerts. You must create a storage pool to enable **All Storage Pools** to appear on the **Storage** tab.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 261.

See [“Alerts and notifications in Backup Exec”](#) on page 303.

To view jobs, job histories, and active alerts for a storage device pool

- 1 On the **Storage** tab, expand **All Storage Pools**.
- 2 Right-click the storage device pool for which you want to view the jobs, and then click **Details**.
- 3 In the left pane, click **Jobs**, **Job Histories**, or **Active Alerts**.

Adding or removing devices in a storage device pool

You can add or remove devices in a storage device pool. Only similar types of storage devices can belong to the same storage device pool. You cannot edit a system-defined storage device pool, such as **Any disk storage**, or **Any disk cartridge storage**.

You must create a storage device pool before **All Storage Pools** appears on the **Storage** tab.

See [“Creating storage device pools”](#) on page 527.

Note: To view or change the default device selection method for a storage device pool, you must use the Backup Exec Management Command Line Interface.

See [“Specifying a default selection method for a device in a storage device pool”](#) on page 530.

To add or remove devices in a storage device pool

- 1 On the **Storage** tab, expand **All Storage Pools**.
- 2 Double-click the storage pool for which you want to add and remove devices.
- 3 In the left pane, click **Properties**.
- 4 Do one or both of the following:

To add a storage device to the pool

Do the following in the order listed:

- Click **Add**.
- Select the devices that you want to add, and then click **OK**.

To remove a storage device from a pool Select the device that you want to remove, and then click **Remove**.

5 Click **Apply**.

Storage operations

This chapter includes the following topics:

- [About storage operation jobs](#)
- [Storage tab overview in Backup Exec](#)
- [Using the Configure Storage wizard](#)
- [Viewing details for multiple storage devices](#)
- [Sending a notification when a scheduled storage operation job completes](#)
- [Scheduling a storage operation job](#)
- [Editing global settings for storage](#)
- [Sharing storage devices](#)
- [Deleting a storage device](#)
- [Changing the state of a storage device to online](#)
- [Renaming a storage device](#)
- [Viewing jobs, job histories, backup sets, and active alerts for storage devices](#)
- [Cataloging a storage device](#)
- [Scanning a storage device](#)
- [Inventorying a storage device](#)
- [Inventorying and cataloging a storage device](#)
- [Pausing and unpausing a storage device](#)
- [Disabling and enabling a storage device](#)

- [Initializing a robotic library](#)
- [Formatting a tape as a WORM tape](#)
- [Retensioning a tape](#)
- [Formatting a tape in a tape drive](#)
- [Ejecting media from a disk cartridge or tape drive](#)
- [Cleaning a robotic library drive](#)
- [Importing media to Backup Exec](#)
- [Exporting media and expired media](#)
- [Locking and unlocking the robotic library's front portal](#)
- [Backup Exec server and storage device states](#)
- [Viewing the disk storage lockdown status](#)

About storage operation jobs

Backup Exec provides the storage operations that help you manage storage devices and media. You can perform most storage operations by right-clicking the storage device, and then selecting the operation. Only the storage operations that are supported for that storage device or media are available on the right-click menu. Not all storage operations are available for all devices.

Backup Exec treats virtual tape libraries and simulated tape libraries as physical robotic libraries. You can identify virtual tape libraries by the label VTL that displays on a library's properties pages. You can identify simulated tape libraries by the label TLS (Tape Library Simulator Utility). The virtual tape libraries and simulated tape libraries do not support all of the storage operations that are available for physical robotic libraries.

You can schedule some storage operations as recurring jobs. You can specify a schedule and a recipient for notification when these jobs run.

You can view all active and scheduled storage operations, and histories of storage operations on the **Job Monitor** tab.

See [“Scheduling a storage operation job”](#) on page 541.

See [“Sending a notification when a scheduled storage operation job completes”](#) on page 541.

Storage tab overview in Backup Exec

On the **Storage** tab, Backup Exec provides overview information for each storage device that you configure. You can view details for each storage device. If you want to view details for all of the storage devices on a server, you can select multiple storage devices.

See [“Viewing details for multiple storage devices”](#) on page 540.

You can customize the columns that appear in the **All Storage** view. Different columns are available from the **Tree** view and from the **List** view. Depending on the view that you select and the storage devices that are available, only some of the columns described in the following table may appear.

See [“How to sort, filter, and copy information on the Backup Exec Administration Console”](#) on page 117.

Table 15-1 All Storage overview

Item	Description
Name	Indicates the name of the storage device. By default, Backup Exec provides a name for the storage device based on the type of storage and an incrementing number, such as Disk storage 0001. You can change the name of the storage device in the storage properties. See “Renaming a storage device” on page 559.
State	Indicates the state of the storage device, such as if it is online, offline, disabled, or if services need to be restarted. See “Backup Exec server and storage device states” on page 584.
Parent Device	Identifies the parent device if there is a relationship between devices. Examples include a relationship between a tape library and a tape drive, or between a storage array and a virtual disk. This column only appears if you select the List view from the Views group at the top of the window.
Backup Exec Server	Identifies the Backup Exec server on which you configured the storage device. If you share the device between multiple Backup Exec servers, the device displays for each Backup Exec server. This column only appears if you select the List view from the Views group at the top of the window.

Table 15-1 All Storage overview (*continued*)

Item	Description
Storage Type	Displays the type of storage that is associated with the device. The storage type can include tape drives, robotic libraries, disk storage, media sets, media vaults, cloud storage, and so on. This column only appears if you select the List view from the Views group at the top of the window.
Active Alerts	Indicates that an event or condition in Backup Exec has occurred for which a message is displayed or a response is required. See “Alerts and notifications in Backup Exec” on page 303.
Storage Trending	Indicates the estimate for the number of days of storage that is left for disk storage and virtual disk storage. See “Storage trending statuses for disk storage and virtual disks” on page 332.

Table 15-1 All Storage overview (*continued*)

Item	Description
Capacity	Displays storage capacity. Backup Exec provides overview information of used and free storage capacity, as well as capacity details for each storage that you configure. Storage capacity information is rolled up for any items that are collapsed under a storage type, such as a robotic library. The information that displays in the Capacity column includes all of the storage capacity of all of the collapsed items. When you expand the items, individual storage capacity information displays. Before capacity information can display for storage, you must inventory and catalog the storage. You can view storage capacity in the following places: ■ On the Storage tab, in the Capacity column. When you hover the mouse over the capacity bar, additional details display in the tool tip. ■ On the Backup and Restore tab, when you specify the storage for a backup job. ■ On the Home tab, in Storage Status. ■ On the Storage tab, when you view properties for disk storage devices. For disk storage and disk cartridge storage, you can set the thresholds for low disk space on the device property pages. See “Editing disk storage properties” on page 338. See “Editing disk cartridge properties” on page 346. See “How Backup Exec catalogs work” on page 254. See “Inventorying a storage device” on page 562.
Total Capacity	Displays the total amount of storage space that is available on the device. For disk storage, this column indicates the size of the volume on which the disk storage is located. For disk cartridge storage, this column indicates the size of the cartridge in the disk cartridge. For tapes, this column indicates the used native capacity and total native capacity of the media.
Used Space	Displays the amount of space that is used as storage, after any compression or deduplication operations have occurred.

Table 15-1 All Storage overview (*continued*)

Item	Description
Backup Data Written	Displays how much raw backup data is backed up, before compression or deduplication have occurred. For example, if you back up 100 MB of data ten times to a deduplication disk folder, the used space is 100 MB but the amount of the backup data written is 1 GB.
Available Space	Displays the difference between Total Capacity and Used Space .
% of Available Space	Displays the difference as a percentage between Total Capacity and Used Space .
Active and Scheduled Jobs	Displays the number of storage operation jobs, backup jobs, and restore jobs that are running on this device. Click the text in this column to see more details of all jobs that are running or that are scheduled to run.
Compression Ratio	Displays the ratio of the uncompressed size of a file over its compressed size.
Average job rate	Displays the average speed of the jobs that run on this device.

Using the Configure Storage wizard

Use the Configure Storage wizard to set up different types of storage to which you can back up data. The Configure Storage wizard creates the storage with the best possible defaults for your environment. However, you can customize all of the device's setting in the device properties.

You can find a list of compatible types of storage devices in the Backup Exec Hardware Compatibility List.

Click **Configure Storage** on the **Storage** tab to start the Configure Storage wizard.

After Backup Exec is installed and the Backup Exec services are started, any storage that is attached to the Backup Exec server is automatically detected. However, you must use the Configure Storage wizard to configure the storage for backups.

Table 15-2 Storage that you can configure in the Configure Storage wizard

Type of storage	Description
Disk-based storage	Storage that remains attached to the server. Types of disk-based storage include the following: ■ Disk storage A location on a locally attached internal hard drive, a USB device, a FireWire device, or a NAS (network-attached storage) device. See “Configuring disk storage” on page 334. ■ Disk cartridge storage Storage that usually remains attached to the server while you remove the media. Disk cartridges use disk cartridge media such as an RDX device, or devices that appear in Windows as removable storage. See “Configuring disk cartridge storage” on page 345. ■ Deduplication disk storage A location on a hard drive that reduces the size of backups by storing only unique data. See “Creating or importing deduplication disk storage” on page 982.
Network storage	Network storage includes the following: ■ NDMP servers Network attached storage (NAS) that supports the Network Data Management Protocol to allow the use of devices that are attached to the servers. See “Features of the NDMP feature” on page 1399. ■ OpenStorage devices Network-attached storage that supports Arctera's OpenStorage technology. See “Configuring an OpenStorage device” on page 419. ■ Cloud storage devices Cloud storage devices are the storage devices that are configured on the cloud hosted by the cloud storage service provider For the list of supported cloud providers, refer to the Backup Exec Hardware Compatibility List. ■ Remote Media Agent for Linux Storage that lets you back up data from remote computers to the storage devices that are directly attached to a Linux server. You can also back up to a simulated tape library on a Linux server. See “About the Remote Media Agent for Linux ” on page 1456.

Table 15-2 Storage that you can configure in the Configure Storage wizard
(continued)

Type of storage	Description
Tape storage	Tape storage includes the following: ■ Stand-alone tape drives Storage that uses a tape cartridge for reading and writing data. See “Support for tape drives and robotic libraries” on page 472. ■ Robotic libraries Storage that contains tape drives, slots, and an automated method for loading tapes. See “Robotic libraries in Backup Exec” on page 478. ■ Barcode rules See “Configuring barcode rules for a robotic library ” on page 480.
Storage pools	Storage pools include the following: ■ Storage device pools ■ Managed Backup Exec server pools See “Creating storage device pools” on page 527.
Media sets and vaults	Media sets and vaults are for tape media only, and include the following: ■ Append period ■ Overwrite protection period ■ Vaulting rules You can also run wizards to update media vaults. See “Managing tapes” on page 491.

See [“About storage operation jobs”](#) on page 534.

Viewing details for multiple storage devices

You can select multiple storage devices for which you want to view all jobs, job histories, and active alerts. Viewing details for multiple storage devices lets you see all of the activity for a specific Backup Exec server.

See [“Storage tab overview in Backup Exec”](#) on page 535.

To view details for multiple storage devices

- 1 On the **Storage** tab, Shift + click or Ctrl + click the storage devices, and then right-click one of the selected storage devices.
- 2 Click **Details**.
- 3 In the left pane, click **Jobs**, **Job History**, or **Active Alerts**.

Sending a notification when a scheduled storage operation job completes

You can assign recipients to be notified when a scheduled storage operation job completes. Recipients must be set up before you can set up notification.

See [“About storage operation jobs”](#) on page 534.

To send a notification when a scheduled storage operation job completes

- 1 Create a new scheduled storage operation job or edit an existing one.
See [“About storage operation jobs”](#) on page 534.
- 2 On the storage operation job dialog box, in the left pane, click **Notification**.
- 3 Select the check box for each recipient that you want to notify when each type of storage operation job completes.
- 4 You can continue selecting other options, or click **OK**.

See [“Scheduling a storage operation job”](#) on page 541.

Scheduling a storage operation job

When you schedule a storage operation job, you can configure the time and frequency that you want to run the job.

See [“About storage operation jobs”](#) on page 534.

To schedule a storage operation job

- 1 On the **Storage** tab, right-click the device for which you want to schedule a storage operation job.
If the storage operation can be scheduled, a small arrow appears next to the operation name.
- 2 Click the storage operation, and then click **Schedule**.
- 3 On the storage operation job dialog box, in the left pane, click **Schedule**.

4 Select any of the following options:

Recurrence	Specify a recurrence schedule for the job.
Hours	Create a recurrence pattern that is measured in hours or minutes. When you select Hours, you can configure the following options: ■ Every X hour/minute Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance. ■ From Designates the starting time for a job to run. ■ Between Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday. The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.
Days	Create a recurrence pattern that is measured in days. When you select Days, you must choose between the following options: ■ Every X day Indicates the number of days between the start time of a job and the start time of the next job instance. ■ Every weekday Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

■ **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

■ **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**

Specifies the number of years between the start time of a job and the start time of the next job instance.

- **On X**

Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.

- **On the X of X**

Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar	View all scheduled jobs on a calendar to check for scheduling conflicts.
Reschedule the job if it does not start x hours after its scheduled start	Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.
Cancel the job if it is still running x hours after its scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.

Create without a schedule

Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

- 5 Click **OK**.

Editing global settings for storage

You can edit the global settings that apply to the robotic libraries, tape media, and disk-based storage that are in your environment.

To edit global settings for storage

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Storage**.
- 3 Select the appropriate options.

Inventory robotic libraries when Backup Exec services start

Enable Backup Exec to inventory all of the slots in a robotic library when Backup Exec services start. Depending on the number of slots and robotic libraries, this process may take a few minutes.

This option is not enabled by default.

Full - protect allocated and imported media

Select this option to prevent Backup Exec from overwriting tape media that are in media sets and the media that are imported from another installation of Backup Exec or from another product.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

This is the safest option to choose because the tape media that is protected cannot be overwritten until one of the following actions occur:

- The overwrite protection period for the media expires.
- You move the media that belongs to an active media set to scratch media.
- You erase, format, or label the media.
- You move imported media to **Scratch Media**.

Partial - protect only allocated media

Select this option to let Backup Exec overwrite tape media that are imported from another installation of Backup Exec or from another product, or scratch media. Media in a media set that has an overwrite protection period that has not expired (allocated media) cannot be overwritten.

It is recommended this option if you want to use media from another installation of Backup Exec or from another product

This option is enabled by default.

Prompt before overwriting imported media

Select this option to be prompted before Backup Exec overwrites tape media that is imported from another installation of Backup Exec or from another product. You must select the option **Partial - protect only allocated media**.

The job cannot run until you respond to this prompt.

None

Select this option to disable the media overwrite protection feature for media in tape drives. With this option, you are responsible for making sure that the media in tape drives are not accidentally overwritten.

When an overwrite job is submitted to a tape drive and the media overwrite protection level is **None**, the media are overwritten.

Note: This option is not recommended because it does not protect data from being overwritten.

Prompt before overwriting allocated or imported media

Select this option to be prompted before Backup Exec overwrites allocated or imported media in tape drives. If you selected **None**, it is recommended that you select this option to be prompted before overwriting allocated or imported media.

The job cannot run until you respond to this prompt.

Overwrite scratch media before overwriting recyclable media contained in the targeted media set

Select this option to let Backup Exec overwrite scratch media first in a tape drive when an overwrite job occurs.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 506.

If no scratch media are found in any of the tape drives, Backup Exec overwrites recyclable media in the selected media set.

If no recyclable media are found in the selected media set, Backup Exec searches for recyclable media in any media set.

If no recyclable media are found, Backup Exec automatically searches for other media to overwrite. The media that is overwritten depends on the level of the overwrite protection that you set. If you select this option, more media may be required for the same number of jobs than if you choose to overwrite recyclable media first.

This option affects the order in which Backup Exec overwrites media. If you choose to overwrite scratch media first, the recyclable media may be preserved longer for possible recovery.

This option is enabled by default.

**Overwrite recyclable media contained
in the targeted media set before
overwriting scratch media**

Select this option to let Backup Exec overwrite recyclable media in a tape drive in the selected media set first when an overwrite job occurs.

If no recyclable media are found in any of the tape drives, Backup Exec overwrites scratch media.

If no recyclable media or scratch media are found, Backup Exec searches for media to overwrite. The media that is overwritten depends on the level of the overwrite protection that you set.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 506.

If you choose to overwrite recyclable media in the selected media set first, the same media is re-used more frequently than if you choose to overwrite scratch media first.

Limit Backup Exec to read-only operations on a disk-based storage device if it has been detached for

Select this option to prevent Backup Exec from reclaiming disk space from expired backup sets on any disk-based storage device that is attached after being absent for a number of days. Backup jobs that you send to this device fail. You can change this setting per disk-based storage device in the device's properties.

Before you disable this setting, you may want to view the expiration dates of the backup sets and decide if you want to keep some backup sets longer. You can change the expiration date of a backup set, or you can retain a backup set indefinitely.

See [“Editing disk storage properties”](#) on page 338.

See [“Backup sets ”](#) on page 359.

See [“How to restore data from a reattached or reinserted disk-based storage device”](#) on page 344.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Number of days

Specify how long before Backup Exec is limited to read-only operations on a disk-based storage device when you reattach it.

The default setting is 14 days.

Limit Backup Exec to read-only operations on a disk cartridge if it has not been inserted for

Prevents Backup Exec from reclaiming disk space from expired backup sets on any disk cartridge that is inserted after being absent for a number of days. Backup jobs that you send to this device fail. You can change this setting per disk cartridge in the device's properties.

Before you disable this setting, you may want to view the expiration dates of the backup sets and decide if you want to keep some backup sets longer. You can change the expiration date of a backup set, or you can retain a backup set indefinitely.

See [“Editing disk cartridge properties”](#) on page 346.

See [“Backup sets ”](#) on page 359.

See [“How to restore data from a reattached or reinserted disk-based storage device”](#) on page 344.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Number of days

Specify the number of days that a device can be absent from the Backup Exec server, after which Backup Exec is limited to read-only operations on the device when you re-insert it.

The default setting is 30 days.

**Allow Backup Exec to delete all expired
backup sets**

Select this option to let Backup Exec delete the last full, incremental, and differential backup sets that are necessary to restore a server if the backup sets have expired. By default, Backup Exec keeps the most recent backup sets that are necessary to restore a server, even if the backup sets expire. If you allow Backup Exec to delete all of the expired backup sets for a server, you may not be able to restore that server.

This option is useful if you do not want to keep data after a period of time, usually several years.

However, when you enable this option, you can lose backup sets in the following conditions:

- When the length of time that the backup data or backup sets is kept is less than the frequency of the backup. That is, the backup sets from the last full backup job expire before the next full backup runs. Ensure that when you create jobs, the backup data is kept longer than the amount of time between full backups.
- When the backup job fails or is missed, and is not rerun before the backup sets expire. Monitor any failed or missed jobs, and ensure that you rerun them before the backup sets from the previous full backup expire.

This option is not selected by default.

Note: In a Central Admin Server feature (CAS) environment, this option is only available on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on all of the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers

in a CAS environment.

See [“Backup sets ”](#) on page 359.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

4 Click **OK**.

Sharing storage devices

In environments in which there is more than one Backup Exec server, those servers can share storage devices. For example, multiple Backup Exec servers in a CAS environment can share storage devices. In these environments, Backup Exec maintains a database of the shared storage device. Otherwise, the backup data that one server submits to the storage device can overwrite the data that another server submits.

Note: The Enterprise Server feature must be installed before you can share storage devices between Backup Exec servers.

Backup Exec servers can share the following types of storage:

- Storage that is attached to an NDMP server
- Deduplication disk storage
- OpenStorage devices
- Cloud storage devices
- Virtual disks
- Disk storage
- Remote Media Agents
- The Backup Exec agents that are configured to send data directly to storage

For disk storage devices and virtual disk, you must specify a UNC path by which the Backup Exec servers can access the storage device. Disk cartridges cannot be shared.

When you share a storage device, you can select which Backup Exec servers can access the storage device. The Backup Exec server from which you added the storage device is automatically enabled to share the storage device. However,

you can remove the sharing capability from that Backup Exec server at any time. For example, if you add a storage device to a central administration server, then that server can use the storage device. However, if your environment does not allow the central administration server to operate as a managed Backup Exec server, then you can remove the sharing capability from the central administration server.

If you have multiple Backup Exec servers and multiple types of storage in your environment, you can select a Backup Exec server and manage the storage for it.

To share a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to share.
- 2 Click **Share**.
- 3 To share a disk storage device or a virtual disk, enter a UNC path by which the servers can access the storage device that you want to share.
- 4 Check the Backup Exec servers or managed Backup Exec servers that you want to share this storage device.
- 5 Click **OK**.

Deleting a storage device

You can delete a storage device from the Backup Exec database. If the storage device is a legacy backup-to-disk folder, a disk storage device, or a deduplication disk storage device, or a cloud-based storage device, then Backup Exec prompts you to delete the backup sets from the Administration Console view. You can no longer view or select those backup sets from the Administration Console. However, the backup sets remain on the storage device. You must run inventory and catalog operations on the storage device before you can restore from it.

You may want to delete the backup sets from the Administration Console if you move a storage device to another Backup Exec installation. However, if the move is only temporary, then you probably do not want to delete the backup sets. Keeping the backup sets lets you avoid running the inventory and catalog operations on the device when you move it back. You should also keep the backup sets if you plan to recreate the storage device.

You can also use Windows Explorer to navigate to a legacy backup-to-disk folder or to a disk storage, and then delete it. If you use this method, then you cannot recreate the storage in Backup Exec.

To delete a storage device

- 1 On the **Storage** tab, right-click the device that you want to delete, and then click **Disable**.
- 2 Right-click the device again, and then click **Delete**.
- 3 When you are prompted to delete the storage device, click **Yes**.

See [“Disabling and enabling a storage device”](#) on page 564.

See [“Backup sets ”](#) on page 359.

Changing the state of a storage device to online

Usually, when a device goes offline, an alert appears. The alert message provides a specific reason why the device is offline.

The alert may include a link to the Knowledge Base for more information.

Correct the problem that caused the device to go offline. Then, for tape drives, robotic libraries, and some other types of devices, you must manually change the state of the device to online.

For disk storage, disk cartridge, storage arrays, and virtual disk devices, Backup Exec detects that the device is online within five minutes and automatically changes the state to online.

See [“Troubleshooting hardware-related issues in Backup Exec”](#) on page 851.

To change the state of a storage device to online

- 1 On the **Storage** tab, right-click the storage device that you want to change to online.
- 2 Click **Offline** to clear the check mark.

Renaming a storage device

You can rename a storage device that is in your environment.

You cannot rename system-defined storage device pools, but you can rename any storage device pools that you create.

See [“About storage operation jobs”](#) on page 534.

To rename a storage device

- 1 On the **Storage** tab, double-click the storage device that you want to rename.
- 2 In the storage device properties, in the **Name** field, type the new name.
- 3 Click **Apply**.

Viewing jobs, job histories, backup sets, and active alerts for storage devices

You can view information that is related to a storage device.

See [“About storage operation jobs”](#) on page 534.

To view jobs, job histories, backup sets, and active alerts for storage devices

- 1 On the **Storage** tab, double-click the storage device for which you want to view the job history, backup sets, or active alerts.
- 2 In the left pane, click **Jobs**, **Job Histories**, **Backup Sets**, or **Active Alerts**.

See [“About the Job History”](#) on page 274.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 261.

See [“About the Job Monitor”](#) on page 263.

See [“Alerts and notifications in Backup Exec”](#) on page 303.

See [“Backup sets ”](#) on page 359.

Cataloging a storage device

You can run a catalog operation to do the following:

- Log the contents of a media that was created by another installation of Backup Exec.
- Create a new catalog on the local hard drive if the catalog for the storage device no longer exists.

Before you can restore or verify data on a storage device, a catalog for that device must exist. If Backup Exec has not used this storage device before, you must run an **Inventory and Catalog** storage operation on the device first.

Note: If a media password was used from a previous release of Backup Exec, Backup Exec catalogs the media as if it is not password-protected. It is recommended that you encrypt data instead.

See [“Using encryption with Backup Exec”](#) on page 729.

See [“Inventorying and cataloging a storage device”](#) on page 563.

See [“How Backup Exec catalogs work”](#) on page 254.

To catalog storage

- 1 On the **Storage** tab, right-click the storage device for which you want to create a catalog.
- 2 Click **Catalog**.
- 3 On the catalog dialog box, click **General**, and then enter a name for the job.
- 4 Click **OK**.
- 5 (Optional) View the job log or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Scanning a storage device

The scan operation gets information about the media that are in the slots, including barcode information if it is available. Then, the scan operation updates the Backup Exec database with the latest information about where the media are located. When you change magazines or insert new media in a magazine in a robotic library, use the scan operation to update the slot information.

The scan job log reports the barcoded media that are in the drives and portals. If the robotic library is busy, the scan job log may not be able to read the drive and portal status. If the drives and portals can't be read, then none are displayed in the job log. For best results, run the scan when the robotic library is idle.

See [“About storage operation jobs”](#) on page 534.

To scan a storage device now

- 1 On the **Storage** tab, right-click the robotic library or slot that you want to scan.
- 2 Click **Scan**, and then click **Scan now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a scan operation for a storage device

- 1 On the **Storage** tab, right-click the robotic library or slot that you want to scan.
- 2 Click **Scan**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 541.

- 5 (Optional) View the scan job log to see which barcoded media are in the slots, drives, and portals of a robotic library, or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Inventorying a storage device

You can run an inventory operation to have Backup Exec read a storage device and update the Backup Exec database with information about the media that is on that device.

For robotic libraries, you can inventory all of the slots in the robotic library when you change tapes. You can also select specific slots to inventory. You are not required to re-inventory slots when you add the tapes that Backup Exec requests. For example, if the data that you want to restore is on a tape that is not in the robotic library, you are prompted to insert the correct tape for the restore operation. In this case, you are not required to re-inventory the slot where the tape is inserted. When you add or remove a tape that Backup Exec does not request, you should run an inventory operation on the changed slots. You can select specific slots to inventory. If you swap tapes often, you may want to run an inventory operation on the robotic library magazine each time that you restart the Backup Exec services.

For tape drives, you can run an inventory operation to mount the media in tape drives and to read the media label. If you change the media that is in a drive, run an inventory operation so that the current media's label appears in the properties. Otherwise, the previous media continues to appear in the properties. There may be a delay as the media is mounted and inventoried in a robotic library.

To inventory a storage device now

- 1 On the **Storage** tab, right-click the storage device that you want to inventory.
- 2 Click **Inventory** and then click **Inventory now** again.

The inventory operation runs. You can view the job log or click the **Job Monitor** tab for details about the job.

To schedule an inventory job for a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to inventory.
- 2 Click **Inventory**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 541.

- 5 Click **OK**.
- 6 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

See [“Inventorying robotic libraries when Backup Exec services start”](#) on page 480.

Inventorying and cataloging a storage device

You can run the inventory and the catalog operations together on a storage device, if the device supports both operations.

See [“About storage operation jobs”](#) on page 534.

To inventory and catalog a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to inventory and catalog.
- 2 Click **Inventory and Catalog**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Pausing and unpausing a storage device

You can pause a storage device to prevent scheduled jobs and new jobs from running on the storage while you perform maintenance activities. Active jobs are not affected if they start before the storage device is paused.

See [“About storage operation jobs”](#) on page 534.

To pause and unpaue a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to pause or unpaue.
- 2 Do one of the following:
 - To pause the storage device, click **Pause**.
 - To unpaue the storage device, right-click it, and then click **Pause** to clear the check mark.

Disabling and enabling a storage device

You can disable a storage device to prevent new jobs from running on it. Backup Exec does not discover disabled NDMP storage devices when the Backup Exec services start.

To disable and enable a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to disable or enable.
- 2 Do one of the following:
 - To disable the storage device, click **Disable**.
 - To enable the storage device, right-click it, and then click **Disable** to clear the check mark.

Initializing a robotic library

You can initialize the robotic library, which sends a startup command to the library.

To initialize a robotic library

- 1 On the **Storage** tab, right-click the robotic library that you want to initialize.
- 2 Click **Initialize**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Formatting a tape as a WORM tape

You can convert a tape to a write-once, read-many (WORM) tape if the tape drive supports the operation. DLT tape drives support the **Format WORM** operation.

To format a tape as a WORM tape

- 1 On the **Storage** tab, right-click the tape drive that contains the tape that you want to convert to a WORM tape.
- 2 Click **Format WORM**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“How WORM media is used in Backup Exec”](#) on page 512.

Retensioning a tape

Before you run a backup job to a tape, you can run the tape in the tape drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. Refer to the documentation that came with your tape drive to see how often to run this operation.

This operation is only available if the tape drive supports retensioning.

See [“About storage operation jobs”](#) on page 534.

To retension a tape

- 1 On the **Storage** tab, do either of the following:
 - Right-click the drive that contains the tape that you want to retension.

- Double-click **Slots**, and then right-click the slot that contains the tape that you want to retension.
- 2 Click **Retension**.
 - 3 (Optional) View the job log or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Formatting a tape in a tape drive

Backup Exec can format the tape in a drive if the drive supports formatting. Formatting a tape may take hours. Most tape drives do not support formatting.

Caution: Formatting erases the tape. All data on the tape is destroyed.

The media label that is displayed was read during the last inventory operation. The media label does not change until another inventory operation occurs. If you change the tape that is in the device but do not inventory the device, the media label that displays may not match the actual media that is in the device.

See [“About storage operation jobs”](#) on page 534.

To format a tape in a tape drive

- 1 On the **Storage** tab, do either of the following:
 - Right-click the tape drive that contains the tape that you want to format.
 - Double-click **Slots**, and then right-click the slot that contains the tape that you want to format.
 - 2 Click **Format**.
 - 3 To format the tape that is displayed, click **Yes**.
 - 4 (Optional) View the job log or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Ejecting media from a disk cartridge or tape drive

Backup Exec can eject the media that is in a disk cartridge or tape drive. Some devices do not support a software-driven media eject. If the media is a tape, the tape is rewound and you may be instructed to manually remove it.

See [“About storage operation jobs”](#) on page 534.

To eject media now from a disk cartridge or tape drive

- 1 On the **Storage** tab, right-click the disk cartridge or tape drive that you want to eject the media from.
- 2 Click **Eject** and then click **Eject now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule an eject operation for a disk cartridge or tape drive

- 1 On the **Storage** tab, right-click the disk cartridge or tape drive that you want to eject the media from.
- 2 Click **Eject**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 541.

- 5 Click **OK**.
- 6 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Cleaning a robotic library drive

You can create and schedule a cleaning job for a robotic library drive.

See [“About storage operation jobs”](#) on page 534.

To create a cleaning job now

- 1 Ensure that you specify the slot that contains the cleaning tape.
See [“Defining a cleaning slot”](#) on page 485.
- 2 Ensure that the cleaning tape is in the defined cleaning slot and is in the same library that contains the drive that you want to clean.
- 3 On the **Storage** tab, right-click the drive that you want to clean, click **Clean**, and then click **Clean now**.
- 4 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a cleaning job

- 1 Ensure that you specify the slot that contains the cleaning tape.
- 2 Ensure that the cleaning tape is in the defined cleaning slot and is in the same library that contains the drive that you want to clean.
- 3 On the **Storage** tab, right-click the drive that you want to clean, click **Clean now**, and then click **Schedule**.
- 4 To send a notification when the job completes, in the left pane, click **Notification** and select the options you want.
- 5 To schedule the job, in the left pane, click **Schedule** and select the options that you want.
See [“Scheduling a storage operation job”](#) on page 541.
- 6 Click **OK**.
- 7 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Importing media to Backup Exec

You can import media to a robotic library to add tapes to Backup Exec, or to import media that is required for a restore job. When you insert media into a robotic library, you must create an import storage operation job. The import storage operation updates the Backup Exec database with the information about the media. Backup Exec associates the media that you import with a system media set.

See [“Default media sets”](#) on page 491.

Note: You should not associate scratch media with a media set that you create. Backup Exec automatically moves the media to the required media set as needed.

Before you import media, note the following:

- If the media does not have a barcode, you must run the **Inventory after import** operation so that the current media's label appears in the properties. You can only select this operation after you select **Import media now**.
- If the robotic library uses a media magazine, ensure that no jobs are currently running. Before you swap the magazine, ensure that all media are ejected from the drive and are back in the magazine slots.

You can select any number of slots to import media to.

The import storage operation supports robotic libraries with portals. When this storage operation job runs, Backup Exec checks the selected slots for media. If media is found, it is exported to the portals. After all of the media is exported, you are prompted to insert new media into the portal so it can be imported. This process continues until all of the requested media have been imported into the robotic library.

You can also run a scan operation to update the slot information when you insert new media in a robotic library. The scan job log reports the barcoded media that are in the drives and portals.

See [“Scanning a storage device”](#) on page 561.

To import media now

- 1 On the **Storage** tab, do one of the following
 - Expand the robotic library, right-click **Slots**, and then click **Import media now**.
 - Right-click the robotic library, and then click **Import media now**.
- 2 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule an import media job

- 1 On the **Storage** tab, do one of the following:
 - Expand the robotic library, right-click **Slots**, and then click **Import media now**.
 - Right-click the robotic library, and then click **Import media now**, and then click **Schedule**.
- 2 In the left pane, click **Storage operations**.

- 3 Click the drop-down menu, and select storage operation that you want to schedule:

Import	Updates the Backup Exec database with information about the media.
Inventory after import	Mounts the media in the drive, reads the media label, and updates the Backup Exec database. This operation is necessary for media that do not have barcodes.

- 4 To send a notification when the job completes, in the left pane, click **Notification** and select the appropriate options:

Recipient name	Show the names of the individual and group recipients.
Recipient type	Indicate Recipient for an individual recipient or Group for a group recipient.
Manage Recipients	Add, edit, or delete recipients.
Properties	View or change the properties of a selected recipient.

- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

Recurrence	Specify a recurrence schedule for the job.

Hours

Create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**

Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance.

- **From**

Designates the starting time for a job to run.

- **Between**

Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.

Days

Create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**

Indicates the number of days between the start time of a job and the start time of the next job instance.

- **Every weekday**

Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

■ **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

■ **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**
Specifies the number of years between the start time of a job and the start time of the next job instance.
- **On X**
Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.
- **On the X of X**
Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar

View all scheduled jobs on a calendar to check for scheduling conflicts.

Reschedule the job if it does not start x hours after its scheduled start

Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.

Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.
Create without a schedule	Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later. If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

6 Click **OK**

7 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Exporting media and expired media

The export media operation supports robotic libraries that have portals. When this operation is run on one or more robotic library slots, the exported media is placed in the portals. If you select more media than there are portals, the robotic library fills as many slots as possible. Then, you are prompted to remove the media from the portal. This process continues until all of the selected media have been removed from the robotic library. You can also export expired media from a robotic library.

The export expired media operation lets you automate media handling in robotic libraries. This operation removes the media that Backup Exec cannot write to. You can then use the **Import after export** operation to add scratch media to the robotic library to prepare for the next backup.

After you export the expired media from the robotic library, the expired media appears in **Offline Tape**. If the media is in a media set that has an applicable vault media rule, then the media appears in the vault location.

You can export cleaning media with the export expired media storage operation. You can include all cleaning media, or all cleaning media that has been used more than a specified number of times.

See [“Importing media to Backup Exec ”](#) on page 568.

Note: Not all storage operations are available for all devices.

To export media or export expired media

- 1 On the **Storage** tab, do one of the following:
 - Expand the robotic library, right-click **Slots**, and then click **Export media**.

- Right-click the robotic library, and then click **Export**.

2 Do one of the following:

To immediately export only the media that Backup Exec cannot write to and place it in the portal

Click **Export expired media now**.

The operation runs. You can view the job history for details about the job.

To immediately export the media to the portal

Click **Export media now**.

The operation runs. You can view the job history for details about the job.

To immediately import media after the export

Click **Import after export**.

The operation runs. You can view the job history for details about the job.

3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

To schedule an export media or export expired media operation

1 On the **Storage** tab, do one of the following:

- Expand the robotic library, right-click **Slots**, and then click **Export media**.
- Right-click the robotic library, and then click **Export**.

2 Click **Schedule**.

3 In the left pane, click **Storage operations**.

- 4 Click the drop-down menu and select one of the following storage operations that you want to schedule:

Export

Places the media into the robotic library's portals.

If you select more media than there are portals, the robotic library fills as many slots as possible. Then you are prompted to remove the media from the portal. This process continues until all of the selected media have been removed from the robotic library.

Import media after export

Adds the scratch media to the robotic library to prepare for the next backup.

Export expired media

Places the expired media into the robotic library's portals.

This operation lets you automate media handling in robotic libraries by removing the media that Backup Exec cannot write to. After you export the expired media from the robotic library, the expired media appears in **Offline Tape**. If the media is in a media set that has an applicable vault media rule, then the media appears in the vault location.

- 5 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 6 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.
- 7 Click **OK**.
- 8 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

Locking and unlocking the robotic library's front portal

By default, the robotic library portal is not locked, even when you run the lock storage operation. Backup Exec Media Servers will not lock robotic libraries when running backup jobs or when a Lock job is run.

If you want to unlock the robotic libraries, contact technical support.

You must create a job to unlock the robotic library's front portal.

To unlock the robotic library's front portal

- 1 On the **Storage** tab, right-click the robotic library that has the front portal that you want to unlock.
- 2 Click **Unlock**, and then click **Unlock now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a job to unlock the robotic library's front portal

- 1 On the **Storage** tab, right-click the robotic library that has the front portal that you want to unlock.
- 2 Click **Unlock**.
- 3 Click **Schedule**.
- 4 To send notification when the job completes, in the left pane, click **Notification** and select any of the following options:

Recipient name	Show the names of the individual and group recipients.
Recipient type	Indicate Recipient for an individual recipient or Group for a group recipient.
Manage Recipients	Add, edit, or delete recipients.
Properties	View or change the properties of a selected recipient.

- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

Recurrence	Specify a recurrence schedule for the job.

Hours

Create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**

Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance.

- **From**

Designates the starting time for a job to run.

- **Between**

Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.

Days

Create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**

Indicates the number of days between the start time of a job and the start time of the next job instance.

- **Every weekday**

Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

■ **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

■ **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

■ **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**
Specifies the number of years between the start time of a job and the start time of the next job instance.
- **On X**
Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.
- **On the X of X**
Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar

View all scheduled jobs on a calendar to check for scheduling conflicts.

Reschedule the job if it does not start x hours after its scheduled start

Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.

Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.
Create without a schedule	Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later. If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

6 Click **OK**.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 560.

Backup Exec server and storage device states

Backup Exec servers and storage devices display a state that indicates their current condition.

Table 15-3 Possible states for Backup Exec servers and storage devices

State	Description
All the Backup Exec services need to be restarted on <Backup Exec server>	The Backup Exec services and the Backup Exec deduplication services must be restarted. See “Starting and stopping Backup Exec services” on page 771.
An error occurred while discovering this device. Cycle the services on <Backup Exec server> to retry device discovery.	The Backup Exec services must be restarted. See “Starting and stopping Backup Exec services” on page 771.
Active	The storage device is in use by a job.
Configuration failed	Configuration has failed for a local disk storage device or virtual disk.
Configuring	A local disk storage device or virtual disk is in the process of configuration.
Disabled	The storage device is disabled and Backup Exec cannot use it. The device is available for other applications.
Disabled; Active	The storage device's status was changed to Disabled while a job was running to the device.
Low disk space; Active	The storage device is in a low disk space condition, but is currently in use by a job.
Low disk space	The storage device has low disk space.
No communication	Communications have stopped between a managed Backup Exec server and a central administration server in a Central Admin Server feature environment. See “What happens when CAS communication thresholds are reached” on page 1357.
Not configurable	The disk cannot be configured because it is in a bad state, or it has failed.
Not configured	The disk is available for configuration but has not yet been configured.

Table 15-3

Possible states for Backup Exec servers and storage devices
(continued)

State	Description
Offline	The storage device is offline. A storage device can appear offline if any of the following actions occur: ■ The device was turned off after Backup Exec started. ■ The device was being used by another application when Backup Exec started. ■ The device is removed from the server. ■ The device reports a critical error. ■ The firmware of the device was updated. Usually, when a device goes offline, an alert appears. The alert message provides a specific reason why the device is offline. The alert may include a link to the Knowledge Base for more information. Correct the problem that caused the device to go offline. Then, for tape drives, robotic libraries, and some other types of devices, you must manually change the state of the device to online. For disk storage, disk cartridge, storage arrays, and virtual disk devices, Backup Exec detects that the device is online within five minutes and automatically changes the state to online See “Changing the state of a storage device to online” on page 559.
Online	The storage device is online.
Paused	The storage device is paused. See “Pausing and unpausing a storage device” on page 564.
Paused; Active	The storage device is paused, but is currently in use by a job.

Table 15-3 Possible states for Backup Exec servers and storage devices
(continued)

State	Description
Stalled	Communications have stalled during communications between a managed Backup Exec server and a central administration server in a Central Admin Server feature environment. See “What happens when CAS communication thresholds are reached” on page 1357.
The Backup Exec deduplication services need to be restarted on <Backup Exec server>	The Backup Exec deduplication services should be restarted. The deduplication services are separate from the Backup Exec services so the Backup Exec services are not affected. See “Starting and stopping Backup Exec services” on page 771.
The Backup Exec services on <Backup Exec server> need to be restarted	The Backup Exec services must be restarted. See “Starting and stopping Backup Exec services” on page 771.
This device has not been discovered correctly. Cycle the services on <Backup Exec server> to retry device discovery.	A state that can occur after you add a new storage device to Backup Exec. You must restart the Backup Exec services so that the device discovery process can run again.
Uninitialized	The device has not been initialized.

Viewing the disk storage lockdown status

The disk storage lockdown setting protects the disk-based backup storage configured with Backup Exec. Access to disk storage is limited only to authorized processes like Backup Exec services. Only Backup Exec is allowed to write to the disk storage. In addition, external processes are not allowed to modify backup data by injecting code into Backup Exec processes. The disk storage lockdown status displays the status of disk-based backup storage configured with Backup Exec.

This setting is enabled by default and is the recommended setting to protect your backup data. You can disable the setting by providing the System Logon Account

credentials. To disable the lockdown, click **Backup Exec Settings > Network And Security > Disk storage lockdown settings > Disable** and enter the System Logon Account credentials.

To view the disk storage lockdown status

- 1 In the Backup Exec status bar, double-click **Disk storage lockdown status**.

The **Disk storage lockdown status details** dialog box is displayed. The **Disk storage lockdown status** displays the status of the disk storage lockdown setting. There can be four states of the disk lockdown.

Enabled	The disk storage lockdown setting is enabled and the disk-based backup storage is protected by Backup Exec.
Disabled	The disk storage lockdown setting is disabled.
Enabled with exception	The disk storage setting is enabled but there are some disk storage for which the lockdown cannot be enabled.
Unavailable	Backup Exec cannot retrieve the status of the disk storage lockdown setting.

- 2 Click **OK**.

See [“Changing network and security options for Backup Exec”](#) on page 718.

Conversion to virtual machines

This chapter includes the following topics:

- [How conversion of physical computers to virtual machines works in Backup Exec](#)
- [Requirements for conversion to virtual machine jobs](#)
- [Converting to a virtual machine simultaneously with a backup job](#)
- [Converting to a virtual machine after a backup job](#)
- [Adding a conversion to virtual machine stage to a backup job](#)
- [Converting to a virtual machine from a point in time](#)
- [Creating a one-time conversion to a virtual machine](#)
- [Setting default options for conversion to virtual machine jobs](#)

How conversion of physical computers to virtual machines works in Backup Exec

Backup Exec provides the ability to convert a physical computer to a virtual machine in the following ways:

- Back up a physical computer and simultaneously convert it to a virtual machine.
- Back up a physical computer and schedule a conversion to a virtual machine to run after the backup job runs.
- Convert existing backup sets to a virtual machine.

- Convert a running physical computer to a virtual machine without running a backup job.

Note: Conversion of a physical Hyper-V host into a virtual machine is not supported. In addition, in a VMware environment, conversion of Windows Server 2012 physical servers that have 4K disks is not supported.

The newly created virtual machine is bootable and is identical to the physical computer from which the virtual machine was converted, with the exception of the network cards and settings. Conversion to a virtual machine enables business continuity for both Hyper-V and VMware environments.

This topic includes the following information:

[Conversion to virtual machine options](#)

[How backup selections are processed during conversion to virtual jobs](#)

[How full, incremental, and differential backups work in conversion to virtual jobs](#)

[Conversion of disks or volumes larger than 2 terabytes](#)

[Notes about conversion to a virtual machine running on a Windows Server 2012 or later Hyper-V host](#)

[Notes about conversion of Exchange servers](#)

Conversion to virtual machine options

You use one of the following options on the **Backup and Restore** tab to set up a conversion to a virtual machine:

Table 16-1 Conversion to virtual machine options

Name of option	Description
Back up to Disk and Simultaneously Convert to Virtual Machine Back up to Deduplication Disk Storage and Simultaneously Convert to Virtual Machine	These options run the conversion simultaneously with the backup job. Because two operations are performed at the same time, this job may take longer to run than a regular backup job. A large backup window is recommended for this option. A conversion from a full backup creates the new virtual machine. Incremental and differential backups update the virtual machine that was created from the full backup. Note: Incremental backups are preferred over differential backups because the differential backups are inefficient for conversion as compared to incremental backups. Although the backup runs simultaneously with the conversion, the backup is the primary job. Therefore, if the backup fails, then the conversion fails also. However, if the conversion fails, the backup continues to run. For a conversion failure, the job is marked as a success with exceptions. In the case of a failed conversion, the conversion process runs again during the next full backup. See “Converting to a virtual machine simultaneously with a backup job” on page 599.

Table 16-1 Conversion to virtual machine options (*continued*)

Name of option	Description
Back Up to Disk and then Convert to Virtual Machine	These options let you schedule the conversion to run after the backup job.
Back Up to Deduplication Disk Storage and then Convert to Virtual Machine	These options require a smaller backup window than the simultaneous conversion options. A conversion from a full backup creates the new virtual machine. Incremental and differential backups update the virtual machine that was created from the full backup. Note: Incremental backups are preferred over differential backups because the differential backups are inefficient for conversion as compared to incremental backups. See “Converting to a virtual machine after a backup job” on page 605.

Table 16-1 Conversion to virtual machine options (continued)

Name of option	Description
Convert to Virtual Machine from Point-in-Time	A conversion to a virtual machine from a point in time converts existing backup sets from a backup job in which all components that are necessary for a virtual machine conversion were selected. When all necessary components are selected for a backup job, Backup Exec identifies that job as Fully selected and the Simplified Disaster Recovery option has a status of ON. The option to convert to a virtual machine from a point in time is useful in a disaster recovery situation in which you want to quickly recover a failed server. The backup sets contain all of the critical components of the server. Additionally, you can select application data or user data to include in the conversion. Note: The option Convert from Point-in-Time becomes available for selection only after you run at least one full backup that includes all critical system components. See “Converting to a virtual machine from a point in time” on page 620.
One-Time Convert to Virtual Machine	This option converts a running physical computer to a virtual machine without a separate backup job. A one-time conversion job can be scheduled to run at a later time, but it cannot be scheduled to run more than one time. Only Full (Copy) backups are supported for this type of conversion. Incremental and differential backups are not supported for one-time conversions. See “Creating a one-time conversion to a virtual machine” on page 622.

Table 16-1 Conversion to virtual machine options (continued)

Name of option	Description
Add Stage	You can add a stage to a backup job to convert to a virtual machine. Two types of stages are available: Convert to Virtual Machine After Backup and Convert to Virtual Machine Simultaneously with Backup. See “Adding a conversion to virtual machine stage to a backup job” on page 613.

Note: Regardless of the option that is used to initiate the conversion, Backup Exec does not power on the virtual machine after creating it.

Backup Exec creates a snapshot of the virtual machine at the end of the conversion process. The snapshot is removed before the next job runs as long as the virtual machine is not powered on and the only snapshot on the virtual machine is the one that Backup Exec created. If you want to start using the virtual machine, you must manually remove the snapshot.

If the converted virtual machine's host fails and you bring the virtual machine online, the existing conversion job continues to run and then fails. In this situation, you must create a new conversion job.

How backup selections are processed during conversion to virtual jobs

When you set up a conversion to virtual job, you select the items to back up before the conversion or simultaneously with the conversion. Backup Exec may automatically exclude or include data in certain situations. The job log lists the items that were excluded from or included in the conversion.

Selections are processed as follows:

- If you exclude a volume from the backup job, then that volume is automatically excluded from the corresponding conversion job.
- If you include an application in the backup job, the disk and volume on which that application resides are automatically included in the conversion job.
- If you exclude a volume from the backup, the disk that is part of the volume is automatically excluded if the volume is the only thing on the disk.

How full, incremental, and differential backups work in conversion to virtual jobs

Conversion-to-virtual-machine jobs create a virtual machine from a full backup. Subsequent incremental and differential backup jobs update the virtual machine that was created by the full backup. Although differential backups are supported for conversion jobs, incremental backups are the preferred method for updating virtual machines. Differential backups are inefficient for conversion when compared to incremental backups.

When using the incremental backup method for conversion to virtual machine jobs, keep in mind the following points:

- Data from volumes is backed up at a file\folder level. Even if only a portion of the file has changed, the entire file is backed up.
- The entire System State is backed up. Incremental backups are supported only for the system file components of System State; the other System State components are backed up as a full backup.
- Block-level backup methods must be selected for SQL. After the initial incremental backup runs for SQL, block-level backups are performed for any subsequent full or incremental backups for which the **Simplified Disaster Recovery** option is **ON**. Conversion will not be performed if a block-level backup method is not selected.
- Full backups are always performed for Active Directory. Neither incremental backups nor differential backups are supported for Active Directory.

Conversion of disks or volumes larger than 2 terabytes

Backup Exec supports conversion of disks or volumes that are larger than 2 terabytes (TB) for VMware hosts and Hyper-V 2012 or later hosts.

Note: If the boot volume or the system volume on the source physical computer is larger than 2 TB and you are converting to a VMware virtual machine, then Backup Exec cannot convert it. A boot volume cannot be split or created as a dynamic spanned volume. This limitation also applies if the boot volume on the source is larger than the destination datastore's maximum supported disk size.

Backup Exec converts all disks and volumes on the destination to simple or spanned dynamic volumes. Even if the source disk is a basic disk, the converted disks on the destination will be dynamic. Since Backup Exec supports only simple and spanned dynamic volumes, if the source volume type is striped, mirror, or raid5, then Backup Exec converts that volume to either simple or spanned on the destination.

The following additional information applies only to VMware:

- Conversion of a volume that is larger than 2 TB is supported in the following situations:
 - If the volume is not a boot or system volume
 - If the volume was created on a dynamic disk
 - If the volume is an MBR disk
- Conversion of GPT disks is not supported, regardless of the size of the volume on the disk. Conversions of GPT disks fail.
- In situations where the VMware source disk size is larger than the destination datastore's maximum supported size, the source disk is split into multiple disks. Backup Exec splits the disk based on the destination datastore's maximum allowed size.

Notes about conversion to a virtual machine running on a Windows Server 2012 or later Hyper-V host

Before you create a job to convert a physical computer to a Windows Server 2012 and later Hyper-V host, review the following information:

- Disk data is stored in vhdx files for conversion of a physical computer to a virtual machine running on a Windows Server 2012 and later Hyper-V host. The vhdx files can have a maximum capacity of 64 TB. The physical computer's disk sector size is maintained during the conversion.
- The conversion of physical computers that have simple GPT disks is supported.
- The conversion of physical computers that have dynamic disks is not supported.
- Storage Spaces and Storage Pools are not supported.
- Conversion to any previous versions of a Windows Hyper-V host is not supported, so those jobs fail. For example, if the physical computer runs Windows Server 2012 with an ReFS volume, conversion to a Windows Server 2012 Hyper-V host is supported.
- If the physical computer runs Windows Server 2012 or later with one or more Windows deduplication volumes, conversion to a Hyper-V host is possible, but it may fail. The converted disk data is not deduplicated. In other words, an unoptimized data transfer is performed. For this reason, the conversion may fail if the amount of unoptimized data is greater than the capacity of the destination volume.

Notes about conversion of Exchange servers

Backup Exec disables Exchange services on a newly-created virtual machine after a conversion-to-virtual job. If Exchange databases are present on a converted server, after a conversion-to-virtual job you must open the Microsoft Services Control Manager and manually restart the following Exchange services:

- MSExchangeDagMgmt
- MSExchangeADTopology
- MSExchangeAntispamUpdate
- MSExchangeDiagnostics
- MSExchangeEdgeSync
- MSExchangeFrontEndTransport
- MSExchangeHM
- MSExchangeMailboxAssistants
- MSExchangeDelivery
- MSExchangeSubmission
- MSExchangeMigrationWorkflow
- MSExchangeMailboxReplication
- MSExchangeRPC
- MSExchangeFastSearch
- HostControllerService
- MSExchangeServiceHost
- MSExchangeThrottling
- MSExchangeTransport
- MSExchangeTransportLogSearch
- MSExchangeUM
- MSExchangeUMCR

See [“Requirements for conversion to virtual machine jobs”](#) on page 597.

Requirements for conversion to virtual machine jobs

Before you use the conversion to virtual machine feature, review the following requirements:

- The option **Simplified Disaster Recovery** must have a status of **ON** on the **Browse** tab of the **Backup Selections** dialog box.

Note: The Agent for VMware and Hyper-V is not required for conversion to virtual machines.

- Only Windows servers are supported.
- Conversion from a duplicate backup set is not supported.
- In a VMware environment, if you convert a physical server to a virtual server and then want to back up the converted server, you must push-install the Agent for Windows to the converted server. Note that you need to install the Agent for Windows on the converted server even if you installed the Agent for Windows on the physical server before you converted it to a virtual server.
- For conversion in a Hyper-V environment, the following additional requirements apply:
 - The Agent for Windows must be installed on the Hyper-V host to which the conversion is sent.
 - Disks larger than 2TB are supported for Hyper-V 2012 or later.
 - Only basic disks are supported. Dynamic disks are not supported for Hyper-V conversions.
 - Before running a conversion job that targets a Hyper-V host, disable the File Server Resource Manager (FSRM). If you do not disable FSRM, then the job may fail with an "out of disk space" error.

Note: Conversion of a physical Hyper-V host into a virtual machine is not supported.

- For specific operating system requirements for conversion to virtual machines, see the Backup Exec Software Compatibility List.

See ["How conversion of physical computers to virtual machines works in Backup Exec"](#) on page 589.

Converting to a virtual machine simultaneously with a backup job

With this type of a conversion, the backup and the conversion run at the same time.

Note: If the backup fails, then the conversion fails also. However, if the conversion fails, the backup continues to run. For a conversion failure, the job is marked as a success with exceptions and the conversion process runs again during the next full backup.

To convert to a virtual machine simultaneously with a backup job

- 1 On the **Backup and Restore** tab, right-click the server that contains the data you want to back up and convert.
- 2 Select **Backup**, and then select **Back up to Disk and Simultaneously Convert to Virtual Machine** or **Back up to Deduplication Disk Storage and Simultaneously Convert to Virtual Machine**, depending on the type of storage device that you want to use.
- 3 Do any of the following:

To change the backup selections	Do the following: ■ In the <Name of Server> box, click Edit, and then select the items to back up. Note: The option Simplified Disaster Recovery must have a status of ON. ■ Click OK.
To change the backup options	Do the following: ■ In the Backup box, click Edit, and then change the backup options as needed. ■ Click OK.

- 4 In the **Conversion to Virtual** box, click **Edit** to set the options for conversion.
- 5 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 6 Configure the conversion options:

If you selected Hyper-V in step 5

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

Note: This option is not available for Hyper-V servers that run on Windows 2016 or later. Such Hyper-V servers install integration services directly on the virtual machine either through a Windows update or a user-initiated download.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk

location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 5

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the

virtual machine bootable.

Note: In a CAS environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

7 Click **OK**.

8 On the **Backup Definition Properties** dialog box, click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 589.

Converting to a virtual machine after a backup job

Backup Exec sets up this type of conversion as a stage that runs after the backup job runs. The backup sets that are created from the backup job are used to create the virtual machine.

To convert to a virtual machine after a backup job

- 1 On the **Backup and Restore** tab, right-click the server that contains the data you want to back up and convert.
- 2 Select **Backup**, and then select **Back Up to Disk and then Convert to Virtual Machine** or **Back Up to Deduplication Disk Storage and then Convert to Virtual Machine**, depending on the type of storage device that you want to use.
- 3 Do any of the following:

To change the backup selections

Do the following:

- In the <Name of Server> box, click **Edit**, and then select the items to back up.

Note: The option **Simplified Disaster Recovery** must have a status of **ON**.

- Click **OK**.

To change the backup options

Do the following:

- In the **Backup** box, click **Edit**, and then change the backup options as needed.
- Click **OK**.

4 In the **Conversion to Virtual** box, click **Edit**.

5 In the left pane, select **Schedule**, and then select one of the following options:

To schedule the conversion to run at a specific time

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Recurrence**, and then click the arrow to set the recurrence pattern.

To create the conversion job without scheduling it

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Create without a schedule**.
 When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

To run the conversion immediately after the backup job completes

Select **Convert to virtual immediately after source task completes**.

- 6 Optional: In the left pane, select **Notification** to notify selected recipients when the job completes.
- 7 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 8 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 9 Configure the conversion options:

If you selected Hyper-V in step 8

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

Note: This option is not available for Hyper-V servers that run on Windows 2016 or later. Such Hyper-V servers install integration services directly on the virtual machine either through a Windows update or a user-initiated download.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk

location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 8

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the

virtual machine bootable.

Note: In a CAS environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

10 Click **OK** to save your selections.

11 On the **Backup Definition Properties** dialog box, click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 589.

Adding a conversion to virtual machine stage to a backup job

You can add a stage to a backup definition to convert a backup to a virtual machine. A conversion to virtual machine job requires that the **Simplified Disaster Recovery** option on the backup selections has a status of **ON**. This status means that all components that are necessary for virtualization are selected. Backup Exec automatically selects the necessary components when you add a stage to convert to a virtual machine.

To add a conversion to virtual machine stage to a backup job

- 1 Create a backup job, or edit an existing job.
- 2 In the **Backup** box, click **Add Stage**.
- 3 Select **Convert to Virtual Machine** to set up a conversion to run after the backup job completes, or select **Convert to Virtual Simultaneously With Backups** to run the conversion at the same time as the backup job.
- 4 In the **Conversion to Virtual** box, click **Edit**.

- 5 If you selected the **Convert to Virtual Machine** option in step 3, do the following. If you selected **Convert to Virtual Simultaneously With Backups** in step 3, skip to step 6.
 - In the left pane, select **Schedule** to schedule the conversion, and then indicate if you want to schedule the job or run it immediately after the backup job completes.
 - (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 6 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 7 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 8 Configure the conversion options:

If you selected Hyper-V in step 7

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

Note: This option is not available for Hyper-V servers that run on Windows 2016 or later. Such Hyper-V servers install integration services directly on the virtual machine either through a Windows update or a user-initiated download.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk

location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 7

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the

virtual machine bootable.

Note: In a CAS environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

9 Click **OK** to save your selections.

10 On the **Backup Definition Properties** dialog box, edit the backup job properties, and then click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 589.

Converting to a virtual machine from a point in time

A conversion to a virtual machine from a point in time converts existing backup sets from a backup job in which the Simplified Disaster Recovery option was enabled. The Simplified Disaster Recovery option enables all of the critical system components for a virtual machine conversion to be included in the backup job.

Note: The option **Convert to virtual machine from point-in-time** becomes available for selection only after you run at least one full backup that includes all critical system components.

The option to convert to a virtual machine from a point in time is useful in a disaster recovery situation in which you want to quickly recover a failed server. The backup sets contain all of the necessary components of the system. Additionally, you can select application data or user data to include in the conversion.

To convert to a virtual machine from a point in time

- 1 On the **Backup and Restore** tab, select the server that contains the backup sets you want to convert.
- 2 In the **Conversions** group, click **Convert to Virtual**, and then click **Convert to Virtual Machine from Point In Time**.
- 3 On the **Options** dialog box, in the **Selected Point-in-Time** box, click **Edit**.
- 4 Select the items that you want to include in the conversion, and then click **OK**.
- 5 On the **Options** dialog box, in the **Convert to Virtual** box, click **Edit**
- 6 In the left pane, select **Schedule** to select when to run the conversion job:

To run the job immediately	Click Run now .
To schedule the conversion to run at a specific time	Click Run on , and then enter the date and time to run the job.
To create the conversion job without scheduling it	Select Create without a schedule . When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

- 7 (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 8 In the left pane, select **Conversion Settings** to set the options for the conversion
- 9 In the **Point in time** field, select the point in time that you want to use for the conversion.
- 10 In the **Name** field, select all of the components for inclusion in the conversion.
- 11 In the **Application data or non-system user data** field, select additional data to include in the conversion.
- 12 Click **OK** to save your selections.
- 13 On the **Options** dialog box, click **OK**.

Creating a one-time conversion to a virtual machine

You can create a one-time conversion job to convert a running physical computer to a virtual machine without a separate backup job. A one-time conversion job can be scheduled to run at a later time, but it cannot be scheduled to run more than one time.

Only Full (Copy) backups are supported for this type of conversion. Incremental and differential backups are not supported for one-time conversions.

To create a one-time conversion to a virtual machine

- 1 On the **Backup and Restore** tab, select the server that you want to convert to a virtual machine.
- 2 In the **Conversions** group, click **Convert to Virtual**, and then click **One-Time Convert to Virtual Machine**.
- 3 On the **One-Time Convert to Virtual Machine Properties** dialog box, in the **Convert to Virtual** box, click **Edit**.
- 4 In the left pane, select **Schedule** to select when to run the conversion job:

To run the job immediately

Click **Run now**.

To schedule the conversion to run at a specific time

Click **Run on**, and then enter the date and time to run the job.

To create the conversion job without scheduling it

Select **Create without a schedule**. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

- 5 (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 6 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 7 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 8 Configure the conversion options:

If you selected Hyper-V in step 7

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

Note: This option is not available for Hyper-V servers that run on Windows 2016 or later. Such Hyper-V servers install integration services directly on the virtual machine either through a Windows update or a user-initiated download.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk

location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 7

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the **ESX/vCenter server name** field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the

virtual machine bootable.

Note: In a CAS environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

9 Click **OK** to save your selections.

10 On the **One-Time Convert to Virtual Machine Properties** dialog box, click **OK**.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 589.

Setting default options for conversion to virtual machine jobs

You can set default options for all conversion to virtual machine jobs. However, you can override the default options for individual jobs.

To set default options for conversion to virtual machine jobs

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select **Convert to Virtual**.
- 3 Select the default schedule options that you want conversion to virtual jobs to use:

To schedule conversion to virtual jobs to run at a specific time

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Recurrence**, and then select the arrow to set the recurrence pattern.

To create conversion to virtual jobs without scheduling them

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Create without a schedule**.
 When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

To run the conversion immediately after the backup job completes

Select **Convert to virtual immediately after source task completes**.

- 4 (Optional) In the left pane, select **Notification**, and then select the recipients who should receive notifications about conversion to virtual jobs.
- 5 In the left pane, select **Conversion Settings**.
- 6 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 7 Configure the conversion options:

If you selected Hyper-V in step 6

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machines, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

Note: This option is not available for Hyper-V servers that run on Windows 2016 or later. Such Hyper-V servers install integration services directly on the virtual machine either through a Windows update or a user-initiated download.

If you selected VMware ESX/vCenter server in step 6

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machines.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CAS environment, this path should be local to the managed Backup Exec server to which the job is targeted.

8 Click **OK**

Configuration and settings

This chapter includes the following topics:

- [Configuring role-based access settings](#)
- [Changing default backup job settings](#)
- [Configuring schedules for backup jobs](#)
- [Configuring storage options for backup jobs](#)
- [Configuring automatic test run jobs for backup jobs](#)
- [Configuring automatic verify operations for backup jobs](#)
- [Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs](#)
- [Configuring Advanced Open File options for backup jobs](#)
- [Configuring checkpoint restart](#)
- [Configuring pre/post commands for backup or restore jobs](#)
- [Configuring parallel streams and job settings for Microsoft 365](#)
- [Configuring file and folder options for backup jobs](#)
- [Setting default schedule options for rule-based jobs and run now jobs](#)
- [Excluding dates from the backup schedule for all backups](#)
- [Removing dates from the list of excluded dates](#)
- [Exporting a list of dates that are excluded from all backups to another server](#)
- [Changing the default preferences](#)
- [Configuring the default setting for backing up multiple servers or applications](#)

- [Configuring database maintenance and security](#)
- [Exporting the Backup Exec Database encryption key](#)
- [Refreshing Backup Exec Database encryption keys](#)
- [Configuring encryption for the connection to the Backup Exec Database](#)
- [Scheduling Backup Exec to check logon accounts](#)
- [Configuring Backup Exec to discover data to back up](#)
- [Adding discovered servers to the list of servers in Backup Exec](#)
- [Backup networks](#)
- [Changing network and security options for Backup Exec](#)
- [Using Backup Exec with firewalls](#)
- [Using encryption with Backup Exec](#)
- [Encryption key management](#)
- [Creating encryption keys](#)
- [Replacing an encryption key](#)
- [Deleting encryption keys](#)
- [Encryption keys and Salt](#)
- [Granular Recovery Technology](#)
- [Setting default Granular Recovery Technology \(GRT\) options](#)
- [DBA-initiated job templates](#)
- [Creating DBA-initiated job templates](#)
- [Editing DBA-initiated job templates](#)
- [Deleting DBA-initiated job templates](#)
- [Backup Exec logon accounts](#)
- [Starting and stopping Backup Exec services](#)
- [Changing the credentials for a service account](#)
- [Changing startup options for Backup Exec services](#)

- [Configuring audit logs](#)
- [Viewing the audit log](#)
- [Removing entries from the audit log](#)
- [Saving an audit log to a text file](#)
- [Copying configuration settings to another Backup Exec server](#)
- [Viewing server properties](#)
- [Configuring default backup settings](#)
- [Changing malware scan options](#)

Configuring role-based access settings

Backup Exec provides the ability to configure role-based access control (RBAC) and gives additional security by controlling the level of access to the Backup Exec console. Use RBAC to assign roles for the Windows users.

By default, role-based access is disabled and can be enabled by providing Owner of System Logon Account credentials. The Owner of the System Logon Account by default has an administrator role and cannot be modified or removed. You must have an administrator role or be Owner of a System Logon Account to be able to manage users within RBAC.

In the CASO environment as well, this feature is disabled by default. You can only enable role-based access control on CAS. After you enable the feature on CAS, it is automatically enabled on MBES. You cannot enable or disable the RBAC feature from MBES servers. You also cannot edit or delete the CAS synchronized users from the MBES servers.

After role-based access control is enabled, only users that are added in the list can log on to Backup Exec or perform operations in BEMCLI.

To configure role-based access settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Role-Based Access Settings**.

The role-based access control feature is disabled by default.

- 2 Click **Enable** and enter the Owner of System Logon Account credentials to enable the feature.

In a CASO environment, ensure that the MBES servers are upgraded to the current version. If any MBES servers are of an earlier version, the list is displayed. It is recommended that you first upgrade all the MBES servers and then enable RBAC in CASO environment

After the credentials are verified, the Owner of System Logon Account has the administrator role. The role cannot be modified or removed.

To disable role-based access control, click **Disable**, enter the Owner of System Logon Account credentials and then click **OK**.

- 3 Do any of the following:
 - Click **Add** to add a new user in the role-based access control list.
The user is displayed in the **Role-Based Access** dialog box along with the roles assigned to each user.
 - Select a user and click **Edit** to update the roles.
 - Select a user and click **Remove** to delete the user from the role-based access control list.
 - Select a user and click **Details** to view additional information about the user and then click **OK**.
 - Click **Synchronize** to synchronize a CAS domain user with all the MBES servers. You can select the user and then click **OK**.
 - Select a user to view the roles that are assigned.
Use the **Roles** section to update the roles.
For more information about the roles and their details, refer to the following section:
See [“RBAC role details”](#) on page 637.

Add a new user in the role-based access control list

You can add new users in the role-based access control list.

To add a new user in the role-based access control list

- 1 On the **Role-Based Access** dialog box, click **Add**.
- 2 On the **Add New User** dialog box, enter the name of a new user account to be added to the role-based access control list.
- 3 Click **Verify** to authenticate the user.

Note: The user that you add must be part of the domain or local Windows administrator group.

- 4 Select the **Connect to the target domain or machine** check box if the new user account for role-based access cannot be verified or if the user is part of a different domain or computer.
- 5 Enter the administrator credentials to connect to the target domain or computer and retrieve the user account details.
- 6 Click **Verify** to authenticate the user.

After the verification is completed, **Windows Group** displays the group that the user is part of.

- 7 Select the roles that you want to assign to the user.

You can also assign multiple roles to the same user.

The following roles can be assigned to a user:

- **Administrator**
- **Storage Administrator**
- **Backup Administrator**
- **Restore Administrator**
- **View Only**

For more information about the roles and their details, refer to the following section:

See ["RBAC role details"](#) on page 637.

- 8 Click **Add**.

The user is added and displayed in the **Role-Based Access** dialog box.

Edit a role in the role-based access control list

You can edit the roles of a user in the role-based access control list.

To edit a role in the role-based access control list

- 1 On the **Role-Based Access** dialog box, select a user and click **Edit**.
- 2 In the **Roles** section, select or remove any roles.
- 3 Click **Save**.

The roles for a user are updated.

Synchronize a CAS user to the MBES servers

In a CASO environment, you can synchronize a CAS domain user to the MBES servers. The synchronization and any changes can only be done from the CAS server and the changes are applied to the MBES servers. You cannot edit or delete the CAS synchronized users from the MBES servers.

To synchronize a CAS user to the MBES servers

- 1 On the **Role-Based Access** dialog box, click **Synchronize** to synchronize a CAS user to the MBES servers.

Note: You can only synchronize a domain user to the MBES servers.

- 2 In the **User Account Synchronization Manager** dialog box, select the users that you want to synchronize across the MBES servers.
- 3 Click **OK**.

On the **Role-Based Access** dialog box, the **Synchronization** column displays the **Enabled** status.

To remove the synchronization for a user, click **Synchronize**, clear the selected check box of the user, and then click **OK**. The synchronization is removed from both CAS and MBES.

Note: Any domain user that is synced on the CAS server cannot be edited or removed from the MBES servers.

RBAC role details

The following table provides details about the different RBAC roles and the operations that you can perform for each role:

Table 17-1 Default RBAC roles

Role	Operations
Administrator	An administrator can perform all operations within Backup Exec. Operations that only an administrator can perform: ■ Manage RBAC users in Backup Exec. ■ Modify media server level settings (Backup Exec settings). ■ Manage licenses. ■ Manage logon accounts. ■ Manage email text notifications. ■ Manage recipients. ■ Manage alerts. ■ Add and remove servers. ■ Add and remove M365 servers. ■ Install Backup Exec agents on guest virtual machines. ■ Delete job history. ■ Export data encryption key. ■ Manage anomalies. ■ Manage reports.
Storage Administrator	Operations that a storage administrator can perform: ■ View Backup Exec settings. Storage administrator must contact administrator to make changes to the settings. ■ Manage all types of storage. ■ Manage the jobs that related to storage, such as inventory and catalog. ■ Perform the Run now operation on any job, put any job on hold or change priority of the job. ■ Respond to alerts. ■ Manage reports.

Table 17-1 Default RBAC roles (*continued*)

Role	Operations
Backup Administrator	Operations that a backup administrator can perform: ■ Manage job defaults. ■ View Backup Exec settings. Backup administrator must contact the administrator to make changes to the settings. ■ Manage backup, verify jobs, and so on except restore and utility jobs. Administrator adds the server, creates the logon account, then storage administrator sets up storage devices, and the backup administrator can create backup jobs. ■ Perform the Run now operation on any job, put any job on hold or change priority of the job. ■ View a list of anomalies. ■ Respond to alerts. ■ Manage reports.
Restore Administrator	Operations that a restore administrator can perform: ■ View Backup Exec settings. Restore administrator must contact the administrator to make changes to the settings. ■ Manage restore jobs and malware scan". Administrator adds the server, creates the logon account, storage administrator runs inventory and catalog if required, and restore administrator can create restore jobs ■ Perform the Run now operation on any job, put any job on hold or change priority of the job. ■ Respond to alerts. ■ Manage reports.
View Only	Operations that a user with view only role can perform: ■ View all running and scheduled jobs. ■ View job history. ■ View alerts. ■ Run the standard reports, run a custom report that is configured, and view the reports that are generated using a scheduled job.

Changing default backup job settings

Backup Exec is preconfigured with default settings for backup jobs. You can change the default settings for your backup jobs. When you create a backup job, the job inherits the default settings that you configure. You can override the default settings for backup jobs when you create them. Backup job settings include storage, security, and file system options for backup jobs, among other things.

You can set unique backup job defaults for the following types of backup jobs:

- Back Up to Deduplication Disk Storage Device
- Back Up to Disk
- Back Up to Tape
- Duplicate to Deduplication Disk Storage Device
- Duplicate to Tape
- Duplicate to Cloud
- Convert to Virtual
- Forever Incremental Backup

Note: Backup Exec displays only the types of backup jobs for which your system is configured. For example, if you do not have a tape drive, you do not see the Back Up to Tape option in the list of backup job types.

To change default backup job settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select the type of backup for which you want to set default options.

For example, if you want to set up the default options for backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select the setting for which you want to configure default options.

Schedule	Select this option to configure default settings for the time and frequency with which you want to run backup jobs. See “Configuring schedules for backup jobs” on page 644.
Storage	Select this option to configure default settings for the storage device that you want to use for backup jobs. See “Configuring storage options for backup jobs” on page 652.
Notification	Select this option to configure Backup Exec to notify specified recipients when backup jobs are completed. Each type of backup job can be configured with different notification recipients. Backup Exec can notify people by email or text message. See “Notification options for jobs” on page 322.
Test Run	Select this option to configure a test job that automatically tests storage capacity, credentials, and media integrity. The test job can help you determine if there are any problems that might keep backup jobs from completing successfully. See “Configuring automatic test run jobs for backup jobs” on page 659.
Verify	Select this option to create a job that automatically verifies whether all of the data was successfully backed up when jobs are completed. A verify job can also help you determine whether the media you use is defective. See “Configuring automatic verify operations for backup jobs” on page 661.

Instant GRT

Select this option to configure Instant GRT or full catalog operations for any GRT-enabled jobs. You can choose to run a full catalog operation immediately after the backup job finishes, schedule the full catalog operation for another time, or run an Instant GRT operation as part of the backup job.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Advanced Open File

Select this option to configure the snapshot settings that Backup Exec uses to process backup jobs. Snapshot technology lets Backup Exec capture any files that are open when a backup job runs.

You can also enable checkpoint restart, which lets you resume interrupted backup jobs.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 669.

See [“Configuring checkpoint restart”](#) on page 672.

Advanced Disk-based Backup

Select this option to configure off-host backup processing for backup jobs.

See [“Setting default backup options for the Advanced Disk-based Backup feature”](#) on page 1385.

Pre/Post Commands

Select this option to configure any commands that you want to run either before backup jobs begin or after backup jobs are completed.

See [“Configuring pre/post commands for backup or restore jobs”](#) on page 674.

Files and Folders

Select this option to configure how Backup Exec processes file system attributes such as junction points and symbolic links.

See [“Configuring file and folder options for backup jobs”](#) on page 683.

Enterprise Vault

Select this option to configure Enterprise Vault options for backup jobs.

See [“Enterprise Vault backup options”](#) on page 1274.

Linux and Macintosh	Select this option to configure options for any Linux or Macintosh computers that are included in backup jobs. See “Linux and Unix backup options” on page 1439.
Microsoft Active Directory	Select this option to configure options for any Microsoft Active Directory data that is included in backup jobs. See “Editing options for Active Directory and ADAM/AD LDS backup jobs” on page 1314.
Microsoft Exchange	Select this option to configure options for any Microsoft Exchange data that is included in backup jobs. See “Setting default backup options for Exchange Server” on page 1190.
Virtual Machines	Select this option to configure options for any virtual machines that are included in backup jobs. See “Setting default backup options for virtual machines” on page 1022. See “Setting default backup options for Hyper-V” on page 1083.
Microsoft SharePoint	Select this option to configure options for any Microsoft SharePoint data that is included in backup jobs. See “Setting default backup options for SharePoint” on page 1211.
Microsoft SQL	Select this option to configure options for any Microsoft SQL data that is included in backup jobs. See “Setting default backup options for SQL Server” on page 1138.
NDMP	Select this option to configure options for any NDMP data that is included in backup jobs. See “NDMP backup options for NDMP servers” on page 1403.
Oracle	Select this option to configure options for any Oracle data that is included in backup jobs. See “Oracle backup options” on page 1247.

Exclusions

Select this option to exclude specific files or specific types of files from backup jobs.

See [“Excluding files from backups”](#) on page 179.

- 4 Select the appropriate options.
- 5 When you are finished configuring default options, click **OK**.

Configuring schedules for backup jobs

Backup Exec lets you configure the time and the frequency for which you want to run jobs. You can run jobs immediately, once on a specific day and time, or more than once according to a schedule. Backup Exec lets you use minutes, hours, days, weeks, months, or years as measurements of time to create a recurring pattern for the schedule. Alternatively, you can select specific days of the month to create a recurring schedule on which jobs should run.

See [“How job scheduling works in Backup Exec”](#) on page 216.

For information about configuring a schedule for a Forever Incremental backup, refer to the following chapter.

See [“Forever Incremental Backup”](#) on page 940.

You can configure default options for schedules, which all your jobs inherit when you create them. Or you can override the default schedule settings when you create jobs.

To configure schedules for jobs

- 1 Do one of the following:

To configure default schedule settings for all backup jobs

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure schedule settings.

To configure a schedule for specific backup jobs

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Schedule**.

- 3 In the **Job template name** field, type the name of the job template for which you want to configure a schedule.

Job templates are the collection of settings that Backup Exec uses to create jobs. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example. When you want to run a backup, Backup Exec combines the job template with the backup selections to create a backup job that runs according to the options that you specified.

The job template name that you enter in this field is used to create the job name.

- 4 In the **Job name** field, type the name of the job for which you want to configure a schedule.

The unique job name helps you to identify backup jobs in Backup Exec. This is only available to set per job, not as a default backup job setting.

- 5 For each job for which you want to configure a schedule, do one of the following:

To configure recurring jobs Complete the following steps:

- Select **Recurrence**.
- Proceed to step 6 to configure the recurrence pattern.

To configure jobs to run immediately without any recurrences

Complete the following steps:

- Select **Run now with no recurring schedule**.
- Proceed to step 9.

Note: This option is only available for full backup jobs.

To configure a job without a schedule Complete the following steps:

- Select **Create without a schedule**.
- Proceed to step 10.

When you use this option, the jobs do not run at the time of creation and they do not have a recurring schedule. The jobs remain unscheduled until you choose to run them. You can use a third-party job automation or task scheduling tool to run the jobs later.

You can run unscheduled backup jobs later by using the **Run Next Backup Now** option or you can manually run the job by using the **Run Now** option.

If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

6 To configure the recurrence pattern, complete the following options:

Hours

Select this option to create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**

Indicate the number of hours or minutes between the start time of a job and the start time of the next job instance.

- You must choose between the following options:

- **From**

Designate the starting time for a job to run.

- **Between**

Restrict the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

The start time and end time can span a maximum of 24 hours. However, they can cross over midnight into the following day.

Days

Select this option to create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

■ Every X day

Indicate the number of days between the start time of a job and the start time of the next job instance.

■ Every weekday

Select this option to enable the job to run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Select this option to create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Select this option to create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

■ **Day X of every X month**

Specify the day on which the job should run. Then indicate the number of months between the start time of a job and the start time of the next job instance.

■ **Every X X of every X month**

Specify the day on which the job should run. Then indicate the number of months between the start time of a job and the start time of the next job instance.

■ **Selected days of the month**

Specify the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. If you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. For example, if you want the job to run on the last Friday of every month, select the check box for Friday in the last row of the grid. Any additional days that you select are added to the monthly recurrence pattern.

■ **Selected dates of the month**

Specify the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. If you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. For example, if you want the job to run on the 1st and 15th of every month, select only those dates in the calendar. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Select this option to create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

■ Every X year

Specify the number of years between the start time of a job and the start time of the next job instance.

■ You must choose between the following options:

■ On X

Specify the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.

■ On the X of X

Specify the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Enter the starting time for the first job in the recurrence pattern.

Starting on

Enter the date on which you want the recurrence pattern to begin.

The date that you enter in this field is the date on which the schedule goes into effect. You can select any date in the past or the future. If you select a date that occurred in the past, Backup Exec calculates the date of the next upcoming job and begins running recurring jobs on that date.

Calendar

Click this option to view all scheduled backup jobs on a calendar to check for scheduling conflicts.

Run initial full backup now in addition to the selected schedule Select this option to run the initial full backup as soon as the job is created without affecting the schedule of future jobs.

Reschedule the job if it does not start X hours after its scheduled start time Specify the amount of time past the job's scheduled start time at which you want Backup Exec to change the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.

See ["List of job statuses in Backup Exec"](#) on page 289.

Cancel the job if it is running X hours after its scheduled start time Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to **Canceled, timed out**.

See ["List of job statuses in Backup Exec"](#) on page 289.

- 7 To configure specific dates to include with the recurring job schedule, select the **Include/exclude dates** tab, and then complete the following steps:

- Click **Include dates**.
- Select the dates that you want to include with the recurring job schedule.
- Click **OK**.

The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.

See ["Including a specific date in the schedule for a backup job"](#) on page 217.

- 8 To configure specific dates to exclude from the job schedule, select the **Include/exclude dates** tab, and then complete the following steps:

- Click **Exclude dates**.
- Select the dates that you want to exclude from the recurring job schedule.
- Click **OK**.

The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.

See ["Preventing backup jobs from running on a specific date"](#) on page 218.

- 9 To submit jobs with an on-hold status, select **Submit job on hold**.

You should select this option if you want to submit the job, but you do not want the job to run until a later date. The job runs later when you change the job's hold status.

- 10 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring storage options for backup jobs

Use the storage options to select the storage and media set on which you want the backup job to run. You can configure different storage devices for each backup job. For example, you may select disk storage for a full backup and a storage pool for an incremental backup in the same backup definition.

You can configure storage options as default settings for all backup jobs. If you do not want to use the default storage options for a particular backup job, you can override the default setting when you create the backup job. You do not have to configure default storage options for all backup jobs, however. If you want to configure different storage options for specific backup jobs, you can configure the storage options when you create those backup jobs.

To configure storage options for backup jobs

- 1 Do one of the following:

To configure default storage options for all backup jobs Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure storage options.

To configure storage options for specific backup jobs Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Storage**.

3 Complete the following options as necessary:

Note: Some of these options display only in Central Admin Server feature (CAS) environments.

Priority	Select the priority of access to the storage devices for backup jobs. This is only available to set per job, not as a default backup job setting. See “Changing the priority for a scheduled job” on page 270.
Backup Exec server or Backup Exec server pool	Select whether you want a job to run on devices on a specific managed Backup Exec server or on devices that are on a group of managed Backup Exec servers. This option displays only if you have the Central Admin Server feature installed. This option is an additional filter that lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed Backup Exec servers in a pool named Exchange Backups, select this option. Then select the Exchange Backups Backup Exec server pool.
Storage	Select the storage device to which you want to send backup data. See “Creating storage device pools” on page 527. See “About the Remote Media Agent for Linux ” on page 1456. See “Features and types of disk-based storage and network-based storage” on page 330.

- 4 If you selected to configure an OpenStorage device or a deduplication disk storage device in the **Storage** field, select from the following options:

Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported

Select this option to enable a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device, and to perform client-side deduplication if the device supports it. The Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If client-side deduplication cannot be performed, then either Backup Exec server deduplication or Appliance deduplication is performed.

This option appears if the Deduplication feature is installed and an OpenStorage device or a deduplication disk storage device is selected in the **Storage** field.

See [“How to use client-side deduplication”](#) on page 995.

Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported

Select this option to enable a remote computer to send data through the Backup Exec server to an OpenStorage device or a deduplication disk storage device, and to perform Backup Exec server-side deduplication if it is supported. If the Backup Exec server does not support deduplication, the data is deduplicated on an intelligent disk device, such as PureDisk or a device from a third-party vendor.

This option appears if the Deduplication feature is installed and an OpenStorage device or a deduplication disk storage device is selected in the **Storage** field.

See [“About the Deduplication feature”](#) on page 969.

Enable Retention Lock

When the **Enable Retention Lock** check box is selected, the backup images that are created, cannot be modified or deleted for a specified retention period.

This check box is available only when the selected OpenStorage device and the logical storage unit have WORM support enabled.

Clear the check box if you do not want the backup job to create backup images that are immutable or indelible.

See [“Configuring an OpenStorage device”](#) on page 419.

LSU Indelible interval

The **LSU Indelible interval** option displays the minimum and maximum retention period applicable to the logical storage unit of a WORM enabled OpenStorage device.

See [“Configuring an OpenStorage device”](#) on page 419.

- 5 In the **Keep for** field, enter the amount of time for which you want to keep the backup sets or job history.

If you have selected **Glacier** or **Deep Archive** cloud storage devices in the **Storage** field, when the backup sets are deleted, the storage provider still incurs charges based on their pricing model.

See [“Notes for cloud-based storage devices”](#) on page 401.

If you have selected an OpenStorage device in the **Storage** field, ensure that the period specified in the **Keep For** field lies between the retention lock interval set on the storage server. This interval is also listed in the **LSU Indelible interval** field.

- 6 If you selected to configure a tape device in the **Storage** field, complete the following options as necessary:

Media set

Select the media set to use for the backup job. The media set specifies the overwrite protection period and the append period for the backup data on the media.

If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu.

This option is available only if you selected a tape device in the **Storage** field.

See [“Default media sets”](#) on page 491.

Overwrite media

Select this option to place the backed up data on overwritable media. Ensure that appropriate media is in the storage device that you select.

Appropriate media for an overwrite job include the following:

- Scratch media
- Media for which the overwrite protection period has expired

Allocated or imported media may also be overwritten depending on the media overwrite protection level that is set.

Depending on your configuration, overwritable media is selected from scratch media or recyclable media.

If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.

This option is available only if you selected a tape device in the **Storage** field.

See [“Managing tapes”](#) on page 491.

See [“Media overwrite protection levels for tape media”](#) on page 505.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 506.

Append to media, overwrite if no appendable media is available

Select this option to append the backed up data to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set.

If the append operation fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert prompts you to insert overwritable media.

This option is available only if you selected a tape device in the **Storage** field.

Append to media, terminate job if no appendable media is available

Select this option to append the backed up data to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job.

This option is available only if you selected a tape device in the **Storage** field.

Eject the media after job completes

Select this option to eject the media from the drive or slot when the operation completes. You can also schedule a job to eject media.

This option is available only if you selected a tape device in the **Storage** field.

See [“Ejecting media from a disk cartridge or tape drive”](#) on page 567.

Retension media before backup

Select this option to run the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.

Use Write once, read many (WORM) media

Select this option to use WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.

See [“How WORM media is used in Backup Exec”](#) on page 512.

Export media to vault after job completes

Select this option to logically move the media from the robotic library to the specified media vault.

This operation moves media from robotic library slots into a portal. An alert reminds you to remove the media from the portal or from a slot. If a job requires multiple media, the export media operation starts after the backup job completes, not after each media is filled.

This option is available only if you selected a tape device in the **Storage** field.

See [“Default media vaults”](#) on page 513.

7 In the **Compression** field, select from the following options:

- None

Select this option to copy the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space.

Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.
- Software

Select this option to use STAC software data compression, which compresses the data before it is sent to the storage device.
- Hardware (if available, otherwise none)

Select this option to use hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed.
- Hardware (if available, otherwise software)

Select this option to use hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.

- 8 To configure encryption, complete the following options, and click **OK**.

Encryption type	Select the type of encryption that you want to use, if any. It is recommended to enable encryption to protect backup data from unauthorized access. For Granular Recovery Technology (GRT) enabled backup jobs and NDMP-based backup jobs, if you have GRT-enabled backup jobs, it is recommended to use storage that supports encryption. You can edit the storage settings or save the job without any change. See “Using encryption with Backup Exec” on page 729.
Encryption key	Select the encryption key that you want to use, if you selected to use encryption. Click Manage the encryptions keys to set the encryption. After you create an encryption key, the selected key is set in the job.
Add Keys	Click this option to create a new encryption key to configure encryption for jobs. Do not use this option to replace or delete any existing encryption keys that are associated with the job. This option is available only if you select an encryption type. See “Encryption key management” on page 733.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring automatic test run jobs for backup jobs

Test run jobs attempt to determine if a scheduled backup job could possibly fail when you run it. When you run a test job, no data is backed up. Instead, Backup Exec checks your storage capacity, credentials, and media to find potential errors. If there is an error, the job continues to run until it is completed. The error appears in the job log. You can also configure Backup Exec to send a notification to a designated recipient.

During a test run job, the following things may cause a job to fail:

- Logon credentials are incorrect.

- Storage capacity is not sufficient.
- Tape cartridge media or disk cartridge media is not available.
- Overwritable media is not available for an overwrite job.
- Appendable media is not available for an append job.

A test run job checks the media capacity that is available for the selected job. However, you can check if there is enough available media for multiple test run jobs in the Test Run Results Report.

See [“Test Run Results report”](#) on page 839.

You can manually run a test run job at any time.

See [“Running a test run job manually”](#) on page 226.

You can also configure test run jobs to run automatically before your scheduled backup jobs as a default setting. If you do not want to run a test run job for a particular backup job, you can override the default setting when you create the backup job. You do not have to enable test run jobs as a default for all backup jobs, however. If you want to run test run jobs only for specific backup jobs, you can configure the test run job when you create those backup jobs.

To configure automatic test run jobs for backup jobs

- 1 Do one of the following:

To enable test run jobs as a default for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure test run jobs.

To enable a test run job for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Test Run**.
- 3 Select **Enable test run**.
- 4 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring automatic verify operations for backup jobs

Backup Exec can perform a verify operation to make sure that the media can be read after a backup job has been completed. It is recommended that you verify all backed up data to ensure the integrity of the collection of data and the media on which it resides.

You can manually run a verify operation on a backup set or a job history at any time. You can verify backup sets if you want to verify only the data that was backed up in a specific backup job instance. If you want to verify a backup definition and all of its dependent backup sets, you can verify a job history. For example, if you want to verify a backup definition that used incremental backups, Backup Exec verifies all incrementals dating back to, and including, the last full backup.

See [“Verifying backed up data manually”](#) on page 227.

By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the automatic verify operation to take place at a later time or disable the verify operation altogether. You can change Backup Exec's verify options as part of the default backup settings or for individual backup jobs.

To configure automatic verify operations for backup jobs

1 Do one of the following:

To configure automatic verify operations for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure verify operations.

To configure automatic verify operations for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Verify**.

3 Complete the following options:

At the end of the job	Select this option to run a verify operation automatically when the backup job is completed.
After job finishes, as a separate job	Select this option to create a verify operation and schedule it to run as a separate job when the backup job is completed. You can use the Edit option to configure options for the separate verify job.
As a separate scheduled job	Select this option to create a verify operation and schedule it to run as a separate job at a later time. You can use the Edit option to configure options for the separate verify job.
Do not verify data for this job	Select this option to disable the verify operation for the backup job. Note: This option is selected by default for the cloud-based storage devices.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs

When you back up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored.

See [“How Backup Exec catalogs work”](#) on page 254.

The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. Backup jobs that are enabled for Granular Recovery Technology (GRT) require more time to catalog because of the amount of granular information that they contain.

On the **Instant GRT and Full Catalog Options** dialog box, you can either choose the Instant GRT option or one of the full catalog options.

Note: The Instant GRT and Full Catalog Options are not supported for backup to tape jobs, cloud storage, or any OpenStorage jobs. If you create a GRT-enabled backup to tape job for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware data, the catalog operation runs as part of the backup job.

See [the section called “Full catalog”](#) on page 663.

See [the section called “Instant GRT”](#) on page 664.

See [the section called “Differences between Instant GRT and full catalog”](#) on page 665.

See [“To configure an Instant GRT or a full catalog operation”](#) on page 666.

Full catalog

For GRT-enabled backup jobs, you can delay the catalog operation and run it as a separate operation to have less effect on your backup window. Because the catalog operation runs separately from the backup job, it does not prevent another scheduled backup job from starting on time.

See [“Granular Recovery Technology”](#) on page 739.

When you enable GRT for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware backups, the full catalog operation runs immediately after the backup job by default.

For Exchange and SharePoint agent-based backups, the full catalog operation runs immediately after all full backups. It runs once every 24 hours for all incremental backups and differential backups, even if you schedule more than one GRT-enabled job to run in the 24-hour period.

For Hyper-V and VMware backups, the full catalog operation runs immediately after all full, incremental, and differential backups by default.

You can also configure the full catalog operation to run on a schedule if you do not want it to run immediately after the backup job.

You may want to schedule the full catalog operation to run outside of your backup window so that it does not interfere with your system resources. If you schedule the full catalog operation, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery.

For example, if you schedule incremental backups to run every 11 hours and schedule the full catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)

- Incremental 1 (10:00 P.M.)
- Catalog 1 (Midnight) This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)
- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight) This job catalogs Incremental 5. Incremental 4 is not cataloged.
- Incremental 6 (5:00 A.M.) This backup is not cataloged.

In the example, the full catalog operation runs only for Incremental 5, Incremental 3, and Incremental 1. For such jobs, you can use the Search wizard to search the data or you can quickly browse for individual items that you want to restore. You can perform a granular recovery using Incremental 2, Incremental 4, and Incremental 6 as well; however, it takes slightly longer to browse items because they are not fully cataloged. Backup Exec dynamically displays the granular data by mounting the backup set.

Instant GRT

An Instant GRT operation runs as part of the backup job and collects only the minimum required catalog information. You cannot use the Search wizard to search the backup sets for individual items. When you browse the backup sets for individual items, Backup Exec reads and displays the granular information in the backup sets as you browse for items that you want to restore. Depending on the backup set you are browsing, full, incremental, or differential, it takes a few minutes or longer to browse individual items.

Before you run an Instant GRT operation, review the following requirements:

- In a CAS environment, ensure that the logon accounts used for backups are added to the list of logon accounts on the central administration server and the managed Backup Exec servers.
- The storage that hosts the backup sets must be online when you browse for individual items that you want to restore because Backup Exec mounts the backup sets dynamically. For incremental and differential backup sets, all such related backup sets should also be accessible during restore.

- If a CAS environment, if a Backup Exec server tries to browse the backup sets of another Backup Exec server and if a firewall is configured between them, you must open ports on the servers.
It is recommended that you browse backup sets either from the managed Backup Exec server on which the backup jobs were run or from the central administrative server.
See [“Backup Exec ports”](#) on page 726.
See [“Backup Exec listening ports”](#) on page 728.

Differences between Instant GRT and full catalog

Table 17-2 Differences between Instant GRT and full catalog

Item	Instant GRT	Full catalog
Granular item search using the Search wizard	Not available.	You can search the backup sets for granular data.
Backup sets browse	You can dynamically browse the backup sets to select the individual items that you want to restore.	You can search as well as browse the backup sets to select the individual items that you want to restore.
Delay in catalog job	No delay in the catalog job. It runs as part of the backup job.	You can configure the catalog job to run immediately after the backup job or at a scheduled time.
Catalog time	No separate catalog time because the cataloging happens as part of the backup job. Backup Exec collects only the minimum required catalog information.	The catalog operation runs as a separate job. It consumes time because Backup Exec collects detailed cataloging information for the backup job.
Catalog file size	Smaller file size because only the minimum required data is cataloged.	Large file size because the complete backup set is cataloged.

Table 17-2 Differences between Instant GRT and full catalog (continued)

Item	Instant GRT	Full catalog
Time to browse granular data for granular restore	Slightly longer than the time taken for browsing data when you do a full catalog because Backup Exec dynamically browses the backup set for reading the GRT data when you expand backup sets to restore granular items. The restore time is the same as for a full-cataloged backup set.	Less than the time taken for browsing data when you do an Instant GRT because the GRT information is already available in the detailed catalog collected during the full catalog job.
Device busy time	The device is not busy for a long time because the catalog is not detailed and it also runs as part of the backup job.	The device is busy for a long time. First when the backup happens, and then when the full catalog operation runs as a separate job.

You can configure these options as default settings for all GRT-enabled backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for Instant GRT or full catalog options, however. If you want to configure these options only for specific jobs, you can configure the settings when you create those jobs.

To configure an Instant GRT or a full catalog operation

1 Do one of the following:

- To configure Instant GRT or full catalog options for all backup jobs
- Complete the following steps:
- Click the Backup Exec button, and then select **Configuration and Settings**.
 - Select **Job Defaults**, and then select the type of backup for which you want to configure the Instant GRT or full catalog settings.

To configure Instant GRT or full catalog options for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Instant GRT**.

3 Select one of the following options:

Enable Instant GRT

Select this option if you want to run an Instant GRT operation for GRT-enabled backup jobs.

This option is the default setting for all new GRT-enabled backup jobs configured on a new Backup Exec installation. After upgrade if this option is not set as the default, you can set this option as the default option to take advantage of faster backups.

For existing GRT-enabled backup jobs that protect Exchange, SharePoint, or virtual machines using the virtual-based backup, the default option for existing jobs does not change. The existing jobs retain the default option that was set before the upgrade. If a new device added after the upgrade, the default option is set to Enable Instant GRT.

If you select this option, the catalog operation runs as part of the backup job and collects only the minimum required catalog information.

You cannot use the Search wizard to search the backup sets for granular data. However, you can browse the backup sets. If you want to restore granular data from the backup sets, Backup Exec browses the backup sets for granular data as you browse for items that you want to restore. If you select this option, the time to browse granular data at the time of restore is slightly longer.

Run a full catalog operation as a separate job immediately after the backup job finishes

Select this option to run the full catalog operation immediately after a backup job finishes. The catalog operation runs as a separate job.

For Exchange and SharePoint agent-based backups, the full catalog operation runs immediately after all full backups. It runs once every 24 hours for all incremental backups and differential backups.

For Hyper-V and VMware backups, the full catalog operation runs immediately after all full, incremental, and differential backups.

Note: Before the full catalog operation completes, instead of using the Search wizard, you must browse the backup sets to select the individual items that you want to restore. The Search wizard is available after the full catalog job is complete.

Schedule a full catalog operation as a separate job after the backup finishes

Select this option to run the full catalog operation as a separate, scheduled job. Then select the start time and the days of the week on which you want the full catalog operation to run.

If you schedule the full catalog operation, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery.

Note: Before the full catalog operation completes, instead of using the Search wizard, you must browse the backup sets to select the individual items that you want to restore. The Search wizard is available after the full catalog job is complete.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

See [“How cataloging works with Hyper-V virtual machine backups”](#) on page 1091.

See [“How cataloging works with VMware virtual machine backups”](#) on page 1031.

Configuring Advanced Open File options for backup jobs

Backup Exec's Advanced Open File feature lets you use snapshot technology to capture any files that are open when a backup runs. You can configure Advanced Open File options as default settings for all backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for Advanced Open File options, however. If you want to use Advanced Open File options only for specific jobs, you can configure the settings when you create those jobs.

To configure Advanced Open File options for backup jobs

- 1 Do one of the following:

- To configure default Advanced Open File options for all backup jobs

Complete the following steps:
 - Click the Backup Exec button, and then select **Configuration and Settings**.
 - Select **Job Defaults**, and then select the type of backup for which you want to configure Advanced Open File options.
- To configure Advanced Open File options for specific backup jobs

Complete the following steps:
 - Create a new backup definition or edit an existing backup definition.
 - In the **Backup** box, click **Edit**.
- 2

In the left pane, click **Advanced Open File**.
- 3

Complete the following options:
- Use snapshot technology**

Select this option to enable the use of snapshot technology for backup jobs.
- Snapshot provider**

Select one of the following snapshot providers for jobs:
 - Automatic - Allow VSS to select the snapshot provider.

Select this option to enable VSS to select the best provider for the selected volume.
 - System - Use Microsoft Software Shadow Copy Provider.
 - Hardware - Use technology provided by the hardware manufacturer.

If you select multiple volumes, you must use the same type of provider to snap all of the volumes. You can snap multiple volumes with the same provider or you can use multiple providers, but you cannot use system and hardware providers as part of the same snapshot.

Process logical volumes for backup one at a time

Select this option to enable the backup of multiple volumes in one job, with only one logical volume being snapped at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped one at a time. A volume with mount points to other volumes is considered a logical volume for snapshot purposes. Therefore, that volume and the mount point volumes are snapped together simultaneously.

After the logical volume is snapped and backed up, the snapshot is detected before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

If this option is not selected, then a snapshot for all volumes in the backup job is created simultaneously. All volumes must meet the minimum quiet time.

This option is only available for local volumes.

The Shadow Copy Components snapshots are created using VSS, which is reported in the job log.

Enable checkpoint restart

Select this option to enable the checkpoint restart option. Checkpoint restart enables Backup Exec to automatically restart a job that is interrupted. The job restarts from the point where it was interrupted instead of starting over at the beginning. Backup Exec waits two minutes after the job stops and then attempts to restart the interrupted job one time. If the job cannot be restarted automatically or if checkpoint restart is disabled, you must restart it manually. A manual restart starts the job at the beginning instead of at the point where the job was interrupted.

See [“Configuring checkpoint restart”](#) on page 672.

4 Click OK.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring checkpoint restart

Checkpoint restart is a backup job setting that enables Backup Exec to automatically restart a job that is interrupted. The job restarts from the point where it was interrupted instead of starting over at the beginning. Backup Exec waits two minutes after the job stops and then attempts to restart the interrupted job. If the job cannot be restarted automatically or if checkpoint restart is disabled, you must restart it manually. A manual restart starts the job at the beginning instead of at the point where the job was interrupted.

Note: Checkpoint restart cannot restart a backup job until it has backed up at least 32 MB of data. If a backup job fails before it has backed up at least that much data, you must run it again manually.

You can enable checkpoint restart for any jobs that are configured with the storage option **Overwrite media**. However, if there is no available media to overwrite and the job is interrupted, it may be placed on hold until overwritable media is available. When the media is available, the job restarts from where it was interrupted.

Backup Exec automatically cancels any jobs that run for too long according to the schedule settings that you selected when you created the job. If Backup Exec automatically cancels a job, it is not eligible to be restarted. If you manually cancel a job, Backup Exec does not automatically try to restart it.

Note: If you use the Central Admin Server feature (CAS), any jobs that are restarted run on the same managed Backup Exec server on which the job failed. If the original Backup Exec server is not available, Backup Exec selects a different Backup Exec server on which to run the restarted job.

You can enable or disable checkpoint restart in the **Advanced Open File** options when you create backup jobs or in the backup job defaults.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 669.

This topic includes the following information:

[Technologies supported for checkpoint restart](#)

[Things to consider before you use checkpoint restart](#)

[Changing the default checkpoint restart settings](#)

Technologies supported for checkpoint restart

Checkpoint restart is only supported for NTFS volumes. The only type of snapshot technology that is supported for checkpoint restart is VSS.

Checkpoint restart is not supported for the following:

- FAT volumes
- FAT32 volumes
- UNIX computers
- Cluster Shared Volumes (CSV)
- Application agents
- Incremental or differential backups
- Jobs that use catalogs to determine if a file has been backed up
See [“How Backup Exec determines if a file has been backed up”](#) on page 198.

Things to consider before you use checkpoint restart

You should consider the following things before you use checkpoint restart:

- If the failure occurs in the middle of an append job, the media is no longer appendable. The media is not appendable until it is erased or overwritten, or the retention period expires. When the restart occurs, Backup Exec uses new media. You should select an appropriate media overwrite protection level to ensure that the restart does not overwrite the media that was used before the job failure.
- If the failure occurs during a verify job or a database consistency check job, the job restarts at the beginning.
- Full backups that were interrupted and resumed from the point of failure do not display in the Simplified Disaster Recovery **Recover This Computer** Wizard. However, you can restore these backup sets manually after you make the initial recovery by using the **Recover This Computer** Wizard.
- You can enable the checkpoint restart option for a full backup job that uses the **Delete selected files and folders after successful backup** option. If the job fails and is resumed, the files are not deleted from the source volume after the backup completes.

Changing the default checkpoint restart settings

You can change the default checkpoint restart settings in the error-handling rules settings. You can specify the number of times that you want checkpoint restart to retry a failed job, the interval between restart attempts, and the final job disposition for any jobs that cannot be successfully restarted.

To change default checkpoint restart settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-Handling Rules**.
- 2 Select **Checkpoint Restart**, and then click **Edit**.
- 3 Select **Enable error-handling rule**.
- 4 Select **Retry job**.
- 5 Complete the following fields:

Maximum retries	Enter the maximum number of times that you want Backup Exec to retry a job that fails.
Retry interval	Enter the interval of time in minutes that you want Backup Exec to wait before it attempts to restart a job.

- 6 In the **Final job disposition** group box, select one of the following options:

Place job on hold until error condition has been manually cleared	Select this option to have Backup Exec place the job on hold if the job cannot be completed successfully after the maximum number of retries. The job remains on hold until the error condition has been manually cleared.
Reschedule job for its next scheduled service	Select this option to have Backup Exec reschedule the job for its next scheduled occurrence if the job cannot be completed successfully after the maximum number of retries.

- 7 (Optional) In the **Notes** field, type any additional notes about the error-handling rule.
- 8 Click **OK**.

Configuring pre/post commands for backup or restore jobs

You can configure commands that you want to run before or after all backup and restore jobs. For example, you may want to create a pre-command to shut down a database before a backup job runs. You could also create a post-command that restarts the database after the job is completed.

Conditions that you can set for these commands include the following:

- Run the backup and restore job only if the pre-command is successful
- Run the post-command only if the pre-command is successful
- Run the post-command even if the backup or restore job fails
- Allow Backup Exec to check the return codes (or exit codes) of the pre- and post-commands to determine if the commands completed successfully
If the pre- or post-command returns an exit code of zero, Backup Exec considers the job to have completed successfully. Backup Exec considers any non-zero exit codes to mean that the job encountered an error.

If it is critical that the job does not run if the pre-command fails, configure Backup Exec to check the return codes. Backup Exec uses the return codes to determine if the pre-command failed or completed successfully.

For example, if a pre-command that shuts down a database before a backup is run fails, the database could be corrupted when the backup runs. In this situation, it is critical that the backup job does not run if the pre-command fails.

If Backup Exec is configured to check the return codes and the post-command returns a non-zero code, the job log reports that the post-command failed. You may have also selected to run the job only if the pre-command is successful. Even if both the pre-command and the job run successfully, Backup Exec marks the job as failed if the post-command fails.

For example, the pre-command can run successfully and shut down the database. The backup job can also run successfully. But if the post-command cannot restart the database, Backup Exec marks the job and the post-command as failed in the job log.

If you select the option **On each server backed up** or **On each server restored** to, the pre- and post-command selections apply to each server independently. The pre- and post-commands are run and completed for one server at a time before they are run on the next selected server.

You can configure pre-commands and post-commands as default settings for all jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for pre-commands and post-commands, however. If you want to use pre-commands and post-commands only for specific jobs, you can configure the settings when you create those jobs.

To configure pre/post commands for backup or restore jobs

- 1 Do one of the following:

To configure default pre/post commands for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure default pre/post commands.

To configure pre/post commands for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- On the **Backup Definition Properties** dialog box, click **Edit**.

To configure pre/post commands for specific restore jobs

Complete the following steps:

- Create a new restore job or edit an existing restore job.
- On the Restore wizard, continue till you come to the pre/post commands.

If you are pre/post commands for restore jobs, skip step 2.

2 In the left pane, click **Pre/Post Commands**.

3 Complete the following options:

Type a command to run before the backup runs	Runs a command on the specified server before the job runs. Use local paths, and make sure that the paths exist on each server and are correct.
Or	
Type a command to run before the restore runs	Commands that require user interaction, such as prompts, are not supported.
Run job only if pre-command is successful	Runs the job only if the pre-command is successful. If the pre-command fails, the job does not run, and is marked as failed. If it is critical that the job does not run if the pre-command fails, then select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully. If a non-zero code is returned, Backup Exec interprets it to mean that the pre-command did not run successfully. The job does not run and the job status is marked as Failed.
Type a command to run after the backup runs	Runs a command on the specified server after the job runs. Use local paths, and make sure that the paths exist on each server and are correct.
Or	
Type a command to run after the restore runs	Commands that require user interaction, such as prompts, are not supported.
Run post-command after job verification completes	Runs the post-command after the verification completes, if you configured a verify operation for the job. See “Configuring automatic verify operations for backup jobs” on page 661.

Run post-command only if pre-command is successful	Runs the post-command only if the pre-command is successful. If it is critical that the post-command does not run if the pre-command fails, then select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully. If a non-zero code is returned for the pre-command, Backup Exec interprets it to mean that the pre-command did not run successfully. The post-command does not run. If you also select Run job only if pre-command is successful, and both the pre-command and the job are successful, but the post-command returns a non-zero code, the job log reports both the job and the post-command as failed.
Run post-command even if job fails	Runs the post-command regardless of whether the job is successful or not. If you also select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully and the post-command returns a non-zero code, the job log reports the post-command as failed.
Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully	Lets Backup Exec check the return codes of the pre- and post-commands to determine if they completed successfully. Backup Exec interprets an exit code of zero from either the pre- or post-command to mean that the command completed successfully. Backup Exec interprets a non-zero code to mean that the command ended with an error. After Backup Exec checks the return codes, it continues processing according to the selections you made for running the pre- and post-commands. If this option is not selected, the success of the pre- and post-commands is not determined based on the return code.
On this Backup Exec server	Runs the pre- and post-commands on this Backup Exec server only.

On each server backed up	Runs the pre- and post-commands one time on each server that is backed up or restored.
Or	
On each server restored to	The pre- and post-command selections apply to each server independently. If you select this option, the pre- and post-commands are run and completed for each server before Backup Exec begins processing on the next selected server.
Cancel command if not completed within x minutes	Designates the number of minutes Backup Exec should wait before it cancels a pre- or post-command that did not complete. The default timeout is 30 minutes.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Configuring parallel streams and job settings for Microsoft 365

You can configure the number of parallel streams for Microsoft 365 backups. It helps to improve the speed of backups by downloading the data in parallel. The data is backed up over the internet from a SaaS service. Microsoft 365 entities like Exchange mailbox, OneDrive volume, SharePoint sites, Teams channel, selected in a job are divided among parallel streams. Each stream protects items from their set of entities and parallelly creates a backup set on the target device. Each stream creates one backup set. The combined view of all the backup sets created by a single backup job is displayed as one point-in-time in the restore browse view of Backup Exec.

You can also update job settings related to Microsoft 365 workloads.

For Microsoft 365 backups, ensure the following:

- The minimum number of streams that you set is less than or equal to the number of concurrent operations that you have defined for the storage device.
- The preferred number of streams is set based on the internet bandwidth, CPU, and RAM requirements. More information is available on the Backup Exec knowledge base.

To configure parallel streams and job settings for Microsoft 365

- 1 Do one of the following:

To configure Microsoft 365 settings for all backup jobs	Complete the following steps: 1 Click the Backup Exec button, and then select Configuration and Settings. 2 Select Job Defaults, and then select the type of backup for which you want to configure Microsoft 365 settings.
To configure Microsoft 365 settings for specific backup jobs	Complete the following steps: 1 Create a new backup definition or edit an existing backup definition. 2 On the Backup Definition Properties dialog box, click Edit.
- 2 In the left pane, click **Microsoft 365**.
 - 3 Complete the following options:

Parallel streams

 - **Preferred number of streams to use for backup:**
 Select the number of streams you want to use for backups. The data is retrieved in multiple streams, which enables your Microsoft 365 backups to run faster.
 The default number of streams is 10.
 - **Minimum number of streams required, fail the job if fewer streams are available:**
 Select the minimum number of streams that you want to use for backups. Ensure that the value you set is less than or equal to the number of concurrent operations that you have defined for the storage device.
 The job fails if streams lower than the minimum number are available.
 The minimum default number is 1.

Exchange Online

■ Retrieve Folder Associated Items:

An Exchange mailbox stores settings (such as, rules) associated with the folders or a mailbox, along with the items. The settings are not displayed in Outlook as regular items. Use this option to back up these settings.

By default, this check box is selected.

■ Retrieve items when Message Read property changes:

Messages may be in unread or read state when they are backed up for the first time. By default, in subsequent backups, messages are not backed up again if their unread or read state changes. Select this option if you want this change to be captured in the subsequent backup.

By default, this check box is not selected.

SharePoint Online

■ Exclude System Lists

The backup excludes hidden SharePoint Online lists and catalogs that are created by default from SharePoint.

■ Exclude Document Libraries with User and Group Sites

If User OneDrive and Group files are backed up separately from Users and Groups selection, select this option to only backup remaining User or Group site data. SharePoint Online excludes data that is backed up under a different selection.

By default, these options are not selected.

Teams

■ Exclude System Lists

The backup excludes hidden SharePoint Online lists and catalogs that are created by default from SharePoint.

■ All Posts

All teams conversations from the beginning are backed up.

■ Date Range

Select the date from which you want to back up the teams conversations.

You have the option to back up the most recent data. It also helps to save costs associated with backing up data.

Channel posts are backed up using Microsoft Graph Export APIs. These are metered APIs and require additional costs.

By default, the **All Posts** option is selected.

About preferred number of streams to use for backup

Performance of the jobs that back up cloud resources depends on factors such as, external network bandwidth between the Backup Exec server and the data center, Microsoft 365 API throttling policies and so on.

Microsoft enforces throttling to limit the number of concurrent calls to a service. When a throttling threshold is exceeded, Microsoft limits any further requests from that client for some time. Throttling helps maintain optimal performance and reliability of the Microsoft services if there are many requests. Microsoft APIs have different throttling limits and behavior.

Backup Exec attempts to get the optimum performance considering all the factors and has some in-built mitigation strategies:

- Utilize available network bandwidth effectively.
Microsoft 365 jobs have multi-stream enabled by default. Backup data is downloaded in parallel streams.
- Retry logic in case of failures is driven by network or throttling issues.
- Automatically reduce the requests to avoid failures due to throttling.
- Data is queried in batches.

A backup job can be configured to run using parallel streams. It helps to increase the speed of the backups by downloading the data in parallel. The data backed up over the internet from a SaaS service.

Microsoft 365 entities like Exchange mailbox, OneDrive volume, SharePoint sites, Teams channel, and so on selected in a job are divided among parallel streams. Each stream protects items from their set of entities and parallelly creates a backup set on the target device. Each stream creates one backup set. The combined view of all the backup sets created by a single backup job is displayed as one point-in-time in the restore browse view of Backup Exec.

To optimize internet bandwidth utilization in a Backup Exec Microsoft 365 job, you can change number of streams used and monitor the impact on the next run of the job. The stream count can be changed using **Backup Options > Microsoft 365 > Preferred number of streams to use for backup**.

The following table recommends the stream count that you can use.

Table 17-3

Available Internet bandwidth	Stream count
Below 50 Mbps	2
50 Mbps	3 to 4

Table 17-3 (continued)

Available Internet bandwidth	Stream count
100 Mbps	5 to 7
150 Mbps	8 to 10

For higher bandwidth, the count depends upon the Backup Exec server hardware and the environment. You need to test different values to check what works for your environment. Refer to the Backup Exec software compatibility list for more information.

The value specified in the **Minimum number of streams required** option is used from the concurrency specified for the target deduplication storage device. The additional number of streams specified in the **Preferred number of streams to use for backup** option after reducing the minimum number of streams is programmatically handled by Backup Exec.

For example, the **Minimum number of streams required** is set to **1** and **Preferred number of streams to use for backup** is set to **10** and the concurrency setting on the deduplication storage is set to **4**.

If the job has as multiple entities selected, to use the preferred streams count of 10, the job uses 1 stream from the concurrency configured on the deduplication storage device and 9 streams are handled programmatically. The remaining concurrency on the deduplication storage (3) is now available for other jobs targeted to deduplication storage device.

The backup job log indicates the number of streams used by the backup job.

Configuring file and folder options for backup jobs

You can configure options for how Backup Exec processes file system attributes such as junction points and symbolic links.

File and folder options can be configured as default settings for all backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for files and folders, however. If you want to configure file and folder settings only for specific jobs, you can configure the settings when you create those jobs.

To configure file and folder options for backup jobs

- 1 Do one of the following:

To configure default file and folder options for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure default file and folder settings.

To configure file and folder settings for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Files and Folders**.
- 3 Complete the following options:

Backup method for files

Select one of the following backup methods:

■ **By modified time**

When Backup Exec runs a full backup job or an incremental backup job, the time that the backup job starts is recorded in the Backup Exec Database. Backup Exec adds the time of the backup job to the Backup Exec Database only if the full backup job completes successfully. The next time that you run an incremental backup job or a differential backup job, Backup Exec compares the file system time to the backup time. If the file system time is later than the time that is recorded in the database, the file is backed up. If the file's modified time is older than the previous backup's modified time, that file is not backed up. If the job does not complete successfully, subsequent differential or incremental backup jobs back up all of the data instead of only the data that has changed.

Note: A file's last modified date and timestamp do not change when the file is copied or moved. To ensure that the files are protected, run a full backup after you copy or move files.

When you run an incremental backup job, Backup Exec records a new time in the Backup Exec Database. The database time is not updated for differential backup jobs.

When you select the modified time method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to the previous backup time to determine if the file has changed.

Using modified time lets Backup Exec run more accurate incremental backups or differential backups even if other processes have modified files' archive bits.

■ **Using archive bit**

Backup Exec uses the archive bit from the file system to determine if a file has changed since the last time it was backed up.

When you use the archive bit, Backup Exec turns the archive bit off when a file is backed up. Turning off the archive bit indicates to Backup Exec that

the file has been backed up. If the file changes again before the next backup job, the bit is turned on again. Backup Exec backs up the file in the next backup.

If the next backup job is a full backup job or an incremental backup job, the bit is turned off when the backup job completes. If the next backup job is a differential backup job, the archive bit is left intact.

■ Using catalogs

Backup Exec compares path names, modified time, deleted and renamed files and folders, and other attributes. When you select the catalog method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to previous catalogs to determine if it has changed.

Note: You cannot use Backup Exec catalogs to determine if files were backed up for any differential backups.

The catalog method is only available if the Advanced Disk-based Backup feature (ADBO) is installed.

Note: The off-host backup feature of ADBO does not support the catalog method.

See [“How Backup Exec determines if a file has been backed up”](#) on page 198.

**Enable single instance
backup for NTFS volumes**

Select this option if you want Backup Exec to check the NTFS volume for identical files. If Backup Exec finds multiple copies of a file, it backs up only one instance of that file.

Single instance backup can considerably reduce the storage space that is required for your backups. Many applications automatically generate some files that have identical content. The actual amount of space that you save depends on the number of duplicate files on the volume.

This option displays only if you use the Microsoft Windows Single Instance Store (SIS) feature.

This option is not applicable if the data targeted for backup is on a Windows Server 2016 and later because Microsoft no longer supports the Single Instance Store (SIS) feature.

Warning: If the backup job does not run to completion, the file data may not be included in the backup set. Rerun the backup job until it is successfully completed. If it was an incremental backup, running the job again does not back up the same files. You must run a full or duplicate backup job to ensure that all files are backed up completely.

Back up files and directories by following junction points and mount points

Select this option to back up the information for the junction points and the files and directories to which they are linked. If this check box is not selected, then only the information for the junction points is backed up. The files and directories to which the junction points are linked are not backed up.

Backup Exec does not follow junction points that are automatically created by Microsoft Windows because it can cause the data to be backed up repeatedly.

You cannot select any mounted drives that do not have a drive letter assigned to them. The files and directories to which they are linked are backed up regardless of whether this option is selected.

If the files and directories to which the junction points are linked are also included in the backup selections, then they are backed up twice. They are backed up once during the full file and directory backup job, and again by the junction point.

Warning: If a junction point is linked to a location that encompasses it, then recursion (a situation where data is backed up repeatedly) occurs. Recursion results in an error and a job failure. For example, if c:\junctionpoint is linked to c:\, recursion occurs when Backup Exec attempts to back up c:\junctionpoint, and the backup job fails.

Back up files and directories by following symbolic links

Select this option to back up the information for any symbolic links and the files and directories to which they are linked.

If you do not select this option, only the information for the symbolic links is backed up. The files and directories to which they are linked are not backed up.

If the symbolic link points to files and directories on a remote computer, the files and directories on the remote computer are not backed up.

Back up data in Remote Storage

Select this option to back up any data that has been migrated from primary storage to secondary storage. The data is not recalled to its original location. It is backed up directly to the backup media.

If this option is selected, you should not run a backup of your entire system. Backup Exec has to load the data that has been migrated to secondary storage and additional time is required for migrated data.

If this check box is cleared, only the placeholder that stores the location of the data on secondary storage is backed up, not the data itself.

This option should not be selected if the device used for secondary storage and backups contains only one drive. If there is only one drive, Remote Storage and Backup Exec compete for use of the drive.

Back up open files

Select one of the following options to determine how Backup Exec processes any open files for the backup job.

The options are as follows:

■ **Never**

Backup Exec skips any open files that are encountered during the backup job. A list of any files that were skipped appears in the job log.

■ **If closed within X seconds**

Backup Exec waits the specified time interval for files to close before it skips them and continues the backup job.

If the file does not close during the specified interval, it is skipped. A list of skipped files appears in the job log.

If multiple files are open, Backup Exec waits the specified time interval for each file. Depending on the number of open files, the wait may significantly increase the backup time.

■ **With a lock**

Backup Exec attempts to open any files that are in use. If Backup Exec is able to open a file, the file is locked while it is backed up. Locking the file prevents other processes from writing to it.

Backing up open files is not as effective as closing applications and allowing the files to be backed up in a consistent state.

■ **Without a lock**

Backup Exec attempts to open any files that are in use. If Backup Exec is able to open the file, the file is not locked while it is backed up. Other applications can write data to the file during the backup operation.

Warning: This option allows some files that contain inconsistent data and possibly corrupt data to be backed up.

Backup method

Select the backup method that you want to use to back up files and folders for each backup job in the backup definition.

See [“Backup methods in Backup Exec”](#) on page 189.

Delete selected files and folders after successful backup

Select this option if you want Backup Exec to delete the data you selected to back up after the backup job completes successfully.

Backup Exec backs up the selected data, verifies the backup sets, and then deletes the data from the server. The logon account credentials that you use to run the job must also have the rights to delete a file. Otherwise, the data is backed up, but it is not deleted.

Note: This option is only available when you configure backup jobs. You cannot configure it as a default setting for all backup jobs.

See [“Configuring Backup Exec to automatically delete files after a backup”](#) on page 202.

Preserve tree on back up and delete

Select this option to retain the file system's directory structure for the files that are backed up in a full backup job. This option is available only when you select the **Delete selected files and folders after successful backup** option.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 640.

See [“Backing up data”](#) on page 158.

Setting default schedule options for rule-based jobs and run now jobs

You can configure default schedule options for rule-based and run now jobs. Backup Exec applies the schedule options whenever you change a run now job or a rule-based job into a recurring scheduled job. A rule-based job is a job that is linked to another job. The rule-based job runs when the job to which it is linked is finished. An example of a rule-based job would be a duplicate stage that is configured to run when a full backup job completes. If you change the duplicate stage's schedule settings, Backup Exec uses the default schedule settings for the duplicate stage. You can override the default settings when you edit the newly scheduled job.

To set default schedule options for rule-based jobs and run now jobs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Schedule**.
- 3 In the **Recurrence Pattern** group box, select the default frequency for backup jobs:

To run jobs every X hours or minutes

Click **Hours**, and then enter the frequency in the **Every X hour/minute** field.

Choose between the following options:

■ **From**

Designate the starting time for a job to run.

■ **Between**

Restrict the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

To run jobs every X days

Click **Days**, and then choose between the following options:

■ **Every X day**

Indicate the number of days between the start time of a job and the start time of the next job instance.

■ **Every weekday**

Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

To run jobs every X weeks

Click **Weeks**, and then specify the number of weeks between the start time of a job and the start time of the next job instance in the **Every X week on** field.

Select the days and time at which jobs should run.

To run jobs every X months Click **Months**, and then choose between the following options:

■ **Day X of every X month**

Specify the specific day on which jobs should run and indicate the number of months between the start time of a job and the start time of the next job instance.

■ **Every X X of every X month**

Specify the day on which jobs should run and indicate the number of months between the start time of a job and the start time of the next job instance.

■ **Selected days of the month**

Specify the days of the month on which Backup Exec should run jobs. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

■ **Selected dates of the month**

Specify the dates of the month on which Backup Exec should run jobs. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

To run jobs every X years Click **Years**, and then enter the frequency in the **Every X year** field.

Choose between the following options:

- **On X**
Specify the date on which Backup Exec should run jobs.
- **On the X of X**
Specify the day and month of the year on which Backup Exec should run jobs.

- 4 (Optional) Click **Calendar** to view all scheduled backup jobs on a calendar to check for scheduling conflicts.
- 5 In the **Reschedule the job if it does not start X hours after its scheduled start time** field, specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.
- 6 In the **Cancel the job if it is still running X hours after its scheduled start time** field, specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
- 7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 216.

See [“List of job statuses in Backup Exec”](#) on page 289.

Excluding dates from the backup schedule for all backups

You can exclude specific dates, such as holidays, from your backup schedule. You may want to exclude holidays, for example, to ensure that Backup Exec does not run jobs on those days. You can exclude dates from the schedule for a specific backup job or you can exclude dates for all backup jobs.

When you exclude dates for all backups, any scheduled backup jobs do not run on those dates. All jobs resume running on their normal schedules after the exclude date. You can still create and run backup jobs and restore jobs on excluded dates, as long as they are not scheduled.

You can exclude dates in Backup Exec by selecting or typing dates on the **Exclude Dates** dialog box. Or you can create a text file with a list of dates to exclude and then import the text file.

After you create a list of dates to exclude, you can export a new text file with those dates. Exporting the text file can be useful if you want to copy your exclude dates from one Backup Exec server to another.

See [“Exporting a list of dates that are excluded from all backups to another server”](#) on page 698.

You can exclude dates from all backups in the job default settings or by using the backup calendar. Both features let you exclude dates from all backups. You can only import dates using the default settings. However, you may prefer to use the backup calendar because it gives you a visual representation of all of your scheduled jobs.

This topic includes the following procedures:

[To exclude dates from the backup schedule for all backups](#)

[To exclude dates from all backups using the backup calendar](#)

To exclude dates from the backup schedule for all backups

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Do any of the following:

To manually enter the date

Complete the following steps:

- In the **Select Date** field, type the date that you want to exclude from the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar

Click the date that you want to exclude.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

To import a list of dates

Complete the following steps:

- Click **Browse**.
- Select the text file that contains the exclude dates.
- Click **Open**.
- Click **Import**.

To delete a date from the list of dates to exclude Complete the following steps:

- Select the date or dates that you want to remove from the list.
- Click **Delete**.

4 When you are finished selecting dates, click **OK**.

To exclude dates from all backups using the backup calendar

- 1 On the **Backup and Restore** tab, in the **Backups** group, click **Backup Calendar**.
- 2 Select the date that you want to exclude from the backup schedule.
- 3 Click **Exclude Dates**.
- 4 Click **Add exclude date for all backups**.

Note: To remove the exclusion from a selected date, click **Remove exclude date for all backups**.

5 Click **Close**.

Removing dates from the list of excluded dates

If you no longer want to exclude a date from your backup schedule, you can remove it from the list of excluded dates. When you remove a date from the list of excluded dates, the date becomes part of your regular backup schedule. Any recurring jobs that normally fall on that day are now scheduled to run rather than being skipped.

To remove dates from the list of excluded dates

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Select the date or dates that you want to remove from the list of excluded dates.

You can also remove dates from the list of excluded dates by clicking the excluded dates on the calendar.

- 4 Click **Delete**.
- 5 When you are finished removing dates from the list, click **OK**.

See [“Excluding dates from the backup schedule for all backups”](#) on page 695.

Exporting a list of dates that are excluded from all backups to another server

You can exclude specific dates, such as holidays, from your backup schedule. When you exclude dates, any regularly scheduled backups do not run on those dates. You can create a list of dates to exclude in Backup Exec.

See [“Excluding dates from the backup schedule for all backups”](#) on page 695.

You can import or export a list of dates to exclude as a text file. This may be useful if you want to copy a list of exclude dates from one Backup Exec server to another.

To export a list of exclude dates

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Click **Export**.
- 4 Browse to the location where you want to save the text file.
- 5 Click **Save**.

Changing the default preferences

You can change the settings for how you prefer Backup Exec to display various screens, indicators, and alerts.

To change the default preferences

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Preferences**.

3 Configure any of the following preferences:

Show splash screen at startup	Clear this option to show the Backup Exec Administration Console instead of the splash screen when you start Backup Exec.
Display the cloud-based help. Clear the check box to view the on-premises help.	Displays the cloud-based help for Backup Exec. You can clear this check-box to view the on-premises help for Backup Exec. By default, this check box is selected. To view the cloud-based help, you must have an internet connection. If you do not have an internet connection, Backup Exec automatically switches to the on-premises help. The default browser available on your machine is used for the cloud-based and on-premises help. It is recommended to use Microsoft Edge, Firefox, Chrome, Opera, and Safari browsers to view the cloud-based and on-premises help.
Display progress indicators for backup jobs. This requires additional time to pre-scan resource.	Select this option to display the percentage complete number while a backup job runs. These indicators appear in the Job Activity dialog box, and they let you monitor the progress of the active job. Backups might take longer to complete when this option is selected because the backup sources must be scanned to determine the amount of data to be backed up. Due to the amount of time that is required to scan the backup sources, you should not select this option when you back up remote resources.
Enable auto synchronization of the licenses with the Licensing Portal.	Enables you to use this option to auto synchronize the licenses with the Licensing Portal. Backup Exec periodically checks for a license update in the Licensing Portal. If an update is available, Backup Exec downloads the updated license file. This option is enabled by default.

Allow Backup Exec to report anonymous usage information (No personally identifiable information will be sent).	Select this option to participate in the Backup Exec product improvement program. General Backup Exec usage and statistical information is periodically collected and sent anonymously to help improve the Backup Exec customer experience. Note: Although usage and statistical information is collected, Backup Exec never collects specific user information.
Use alternating row colors	Select this option to display alternating row colors for all lists in Backup Exec. Alternating row colors can make it easier to distinguish between rows.
Re-enable	Select this option to enable any messages that you have disabled.
Enable Security Event Logging	Enables you to log application security events such as logon failures, network connection failure, and job failures. It is recommended to use this setting only when it is required as the events are logged into the Windows Event Log.
Anomaly Detection Status	Enables you to manage the Anomaly Detection feature. By default, the feature is enabled. Click Disable to disable the feature. If you click Disable, you must provide the credentials for the Owner of System Logon Account. Anomaly detection stops once the feature is disabled. You are not required to provide credentials for enabling the feature.

4 Click **OK**.

Configuring the default setting for backing up multiple servers or applications

You can select to back up multiple servers or applications at once using Backup Exec. You can back them up as part of one backup definition or you can back them up individually in separate backup definitions. It may be easier for you to manage backing up multiple servers as part of one backup definition. However, it is easier to troubleshoot job failures if each server has its own backup definition.

Each time you create a backup definition that contains multiple servers or applications, you can select whether you want to create one backup definition or separate backup definitions. You can configure a default scenario for backing up multiple servers or applications so that Backup Exec automatically creates either one backup definition or separate backup definitions.

To configure the default setting for backing up multiple servers or applications

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Backups**.
- 3 Select one of the following:
 - **Create one backup that includes all servers or applications**
 - **Create separate backups for each server or application**
- 4 If you want Backup Exec to prompt you each time you select to back up multiple servers or applications, select **Prompt each time I create a backup for multiple servers**.

If you disable the prompt, Backup Exec automatically uses the preference that you selected in the **When creating backups for multiple servers** field. You can enable the prompt at any time.

See [“Backing up data”](#) on page 158.

Configuring database maintenance and security

The Database Maintenance and Security option lets you manage the Backup Exec Database. Each database maintenance operation is performed independently on each database. The Backup Exec Database maintains a record of the files and data that you have configured.

Database maintenance lets you perform the following:

- Optimize database size.
- Delete expired data.
- Save the contents of the database files.
- Perform a database consistency check.

Backup Exec generates informational alerts at the beginning and at the end of the database maintenance process each time database maintenance is performed. The alerts provide details about the type of maintenance that was performed on each database and the amount of time that the maintenance took to complete. If

the database maintenance process fails, the alert indicates where the failure occurred and the reason for the failure.

You do not have to select all the options; however, each one performs a different process that enables you to protect and maintain your database. Selecting all the options enables you to recover the database quickly and maintain optimal performance.

You can also export the Backup Exec Database encryption key. The Backup Exec Database encryption key is used to secure the Backup Exec Database. The key is required for a number of disaster recovery and migration scenarios. You should export the encryption key to a safe location to ensure that you have it later.

To configure database maintenance and security

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Database Maintenance and Security**.
- 3 To enable database maintenance, select the **Enable Backup Exec database maintenance** option.

4 Configure any of the following options:

Perform database maintenance daily at	Select the time that you want to perform database maintenance. All the maintenance occurs once a day at the time you specify.
Delete aged data	Select this option to delete expired job history, job logs, alert history, and reports from the Backup Exec Database after the specified number of days have passed.
Keep job history for data on media that have current overwrite protection periods	Select this option to keep all job history data for any media to which an overwrite protection period is currently assigned. After a media's overwrite protection period expires, the media's job history data can be deleted.
Keep job history for specified number of days	Select this option to indicate the number of days to keep job history data in the database before it is deleted. Job history data includes summary statistics for a job and details about media, devices, and any backup sets that were used to process the job.
Job logs	Indicate the number of days to keep job logs in the database before they are deleted. Job logs include detailed information about the job.
Alert history	Indicate the number of days to keep alert history data in the database before it is deleted. Alert history data includes property and response information for the alert.
Reports	Indicate the number of days to keep report data in the database before it is deleted. Report data includes property information about any report jobs that were generated. The report itself is not deleted.
Audit logs	Indicate the number of days to keep audit log data in the database before it is deleted. The audit log includes information about any operations that are performed in Backup Exec. See "Configuring audit logs" on page 774.

Perform database consistency check

Select this option to check the logical consistency and physical consistency of the data in the database.

The option is not checked by default. It is recommended that you run a consistency check periodically at a time when there is minimal activity from Backup Exec.

Save contents of database to the Backup Exec data directory

Select this option to save the data that is contained in the database to the Backup Exec data directory so that the database backup file (BEDB.bak) can be backed up.

The dump file is maintained in the data directory until the next database maintenance process is performed and then this file is overwritten. Selecting this option enables you to recover the database in the event of failure.

Optimize database size

Select this option to organize fragmented pages and decrease the size of the physical database to 10 percent above what is actually used.

- 5 To export the database encryption key, complete the following fields:

Note: You should export the encryption key to a safe location to ensure that you have a copy of it for later. You need the encryption key to perform disaster recovery or migrate the Backup Exec server. The key is named with a unique hash value. Backup Exec uses the name to identify the key later.

See [“Exporting the Backup Exec Database encryption key”](#) on page 705.

Path	Type the path of a secure location to which you want to export the Backup Exec Database encryption key.
Remember the export path. By clicking this check box, you consent to let Backup Exec retain and display the export path during import operations.	Select this option to let Backup Exec remember the path to which you exported the database encryption key. If you select this option, Backup Exec can attempt to automatically recover the database encryption key in the event that the key fails. If you do not select this option, you must manually import the database encryption key if it ever fails.
Export	Click this option to export the Backup Exec Database encryption key to the location that you specified in the Path field.

- 6 Click **OK**.

See [“Configuring encryption for the connection to the Backup Exec Database”](#) on page 708.

Exporting the Backup Exec Database encryption key

Backup Exec stores sensitive information in the Backup Exec Database using encryption. When you install or upgrade Backup Exec, it automatically creates a database encryption key. The database encryption key is used to encrypt information such as login account credentials and the keys that are used for encrypted backup jobs, for example. It is stored in the Data folder in the Backup Exec installation directory.

You are required to provide the Backup Exec Database encryption key for each of the following scenarios:

- Performing a manual disaster recovery of a Backup Exec server

- Performing a disaster recovery of a Backup Exec server using Simplified Disaster Recovery (SDR)
- Migrating Backup Exec from one computer to another computer
- Resolving any situations in which the database encryption key on the Backup Exec server is corrupted or goes missing

It is recommended that you export the Backup Exec Database encryption key to a secure location so that you can access it later if it is needed. You should repeat the following procedure on each Backup Exec server in your environment, including the central administration server and each managed Backup Exec server in Central Admin Server feature (CAS) deployments.

Make sure that you export the database encryption key to a location that meets the following criteria:

- The destination is either on a physical volume that is assigned to a drive letter or a network share that is specified by a UNC path (network shares that are mapped to drive letters are not supported)
- The destination has enough disk space
- The destination is accessible from the Backup Exec server
- Backup Exec has permission to write to the destination

To export the Backup Exec Database encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Database Maintenance and Security**.
- 3 In the **Path** field, type the location to which you want to export the encryption key.
- 4 If you want Backup Exec to remember the path to which you exported the database encryption key, select **Remember the export path. By clicking this check box, you consent to let Backup Exec retain and display the export path during import operations.**

If you select this option, Backup Exec can attempt to automatically recover the database encryption key in the event that the key fails. If you do not select this option, you must manually import the database encryption key if it ever fails.

5 Click **Export**.

The key is exported to the location that you specified. The key is named with a unique hash value. Backup Exec uses the name to identify the key later. If you want to export the key to additional locations, repeat steps 3 through 5.

6 Click **OK**.

See [“Configuring database maintenance and security”](#) on page 701.

See [“Refreshing Backup Exec Database encryption keys”](#) on page 707.

See [“Configuring encryption for the connection to the Backup Exec Database”](#) on page 708.

Refreshing Backup Exec Database encryption keys

Backup Exec stores sensitive information in the Backup Exec Database using encryption. A database encryption key is used to encrypt information such as login account credentials and the keys that are used for encrypted backup jobs, for example. The key is stored in the Data folder in the Backup Exec installation directory. It is required for many disaster recovery and migration scenarios.

Backup Exec automatically creates the Backup Exec Database encryption key. However, you may want to refresh the key if the existing key is compromised in any way. You may also be required to change the key if your organization requires that keys or passwords be changed periodically.

Note: You must have a functioning database encryption key to complete the procedure below.

Complete the following procedure to refresh the Backup Exec Database encryption key.

To refresh Backup Exec Database encryption keys

- 1 In Windows, click **Start** and then click **Run**.
- 2 Type **Regedit** and then click **OK**.

Warning: Incorrect use of the Windows registry editor may prevent the operating system from functioning properly. You should take great care when you make changes to the Windows registry. Registry modifications should only be carried out by persons who are experienced in the use of the registry editor application. It is recommended that you make a complete backup of the registry and computer before making any registry changes.

- 3 Locate and right-click the following registry key:
HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\Backup Exec For Windows\Backup Exec\Server\DatabaseEncryptionAction
- 4 Click **Modify**.
- 5 In the **Value data** field, type 2.
- 6 Click **OK**.
- 7 Restart all Backup Exec services.

Backup Exec creates a new Backup Exec Database encryption key. It is recommended that you export the new key to a secure location so that you can access it later if it is needed.

See [“Exporting the Backup Exec Database encryption key”](#) on page 705.

Configuring encryption for the connection to the Backup Exec Database

The Backup Exec Database contains sensitive information about your organization, including user account credentials and backed up data. Securing Microsoft SQL Server's connection to the Backup Exec Database is an important step in protecting your network from outside access. Microsoft recommends that you use SSL encryption any time data that is transmitted between SQL Server and an application travels across a network.

Data transmission between the Backup Exec services and the SQL instance can travel across the network in the following scenarios:

- You configure the Backup Exec Database as a centralized database and it is located on a central administration server in a CAS environment. Data can

also travel across the network in variations of this scenario, for example when you use a managed Backup Exec server or when you use shared storage.

- You use a remote SQL instance for the Backup Exec Database so that the Backup Exec services must access the database across the network.

Backup Exec automatically enables SSL encryption if you use the default, local SQL Express instance called "BKUPEXEC". If you configure Backup Exec to use any other SQL Server instance, you must configure encryption yourself.

SQL Server uses certificates to encrypt data. You can generate your own certificates or you can let SQL Server use an automatically generated, self-signed certificate. By default, Backup Exec uses the self-signed certificates that SQL Server automatically generates. However, It is recommended that you create and use your own certificates for additional security.

Note: Using encryption may affect the performance of communications between SQL Server and the Backup Exec Database. It involves an extra round trip across the network as well as time to encrypt and decrypt the data.

Refer to the Microsoft knowledge base for more information about Secure Sockets Layer (SSL) and encrypting connections to SQL Server.

For information about the best practices to manage the database encryption, refer to the *Backup Exec Best Practices*.

To generate and install certificates for secure SQL connections (optional)

You can use your own certificates or you can let SQL Server use an automatically generated, self-signed certificate. It is recommended that you use your own certificates for improved security. Once you have generated and installed your certificate, you can proceed to configure the secure SQL connection to the Backup Exec Database.

Microsoft has requirements that must be followed when you use your own certificates for SQL Server. Certificates can be either self-signed or issued from a certification authority. Certification authorities can be either a local authority in your organization's domain or a known third-party authority.

For more information about Microsoft's certification requirements, refer to the following Microsoft article:

[Encrypting Connections to SQL Server](#)


Before you configure encryption, you must import the certificates that you want to use into the local certificate store of the computer that hosts the Backup Exec Database.

For more information about importing and installing a certificate on the server, refer to the following Microsoft article:


[How to: Enable Encrypted Connections to the Database Engine \(SQL Server Configuration Manager\)](#)

When you import certificates, you should use the same user account under which the SQL Server service runs:

- If the SQL Server is running under a default computer account such as LocalSystem, NetworkService, or LocalService, then you should use the **Computer account** option when you import the certificate. Selecting to manage certificates for the computer account ensures that the certificate is placed under the Personal store of the default computer account.



- If the SQL Server is running under a specific domain account, you must be logged in using the same domain account to import the certificate. When you log into the Microsoft Management Console, select the **My user account** option. Selecting to manage certificates for the user account ensures that the certificate is placed under the Personal store of the user who is also running the SQL service account.



To configure secure SQL connections to the Backup Exec Database

Backup Exec automatically enables encryption for SQL connections if you use the default, local SQL Express instance called "BKUPEXEC". If you configure Backup Exec to use any other SQL Server instance, you must configure encryption yourself. You should configure the secure connection on the computer on which the SQL instance hosts the Backup Exec Database.

In some Backup Exec environments, you may need to configure the secure connection more than once:

For Central Admin Server feature (CAS) environments	You must configure a secure SQL connection on each computer in the CAS environment, including the central administration server and any managed Backup Exec servers.

Use the SQL Server Configuration manager to edit the properties of the protocols for the server that you want to configure. If you want to configure encryption for the default, local database instance that Backup Exec installs, edit the **Protocols for BKUPEXEC**. Select the certificate that you want to use, if you created a

certificate. Then select whether you want to force encryption for the database connection. When you have finished, restart SQL Server and the Backup Exec services from the Services Manager.

For more information or instructions for configuring encrypted connections for SQL, refer to the Microsoft knowledge base.

Scheduling Backup Exec to check logon accounts

You can schedule Backup Exec to check that the backup sources in jobs can be accessed with the logon accounts that you selected. Checking whether your logon accounts have access to backup sources lets you diagnose and fix any access issues before you run backup jobs. If Backup Exec discovers any backup sources that cannot be accessed with the logon accounts that you selected, it reports the error in an alert.

By default, Backup Exec is scheduled to check logon accounts every day at 2:00 pm.

You can disable the test if you do not want Backup Exec to regularly check logon accounts. You can also reschedule the test so that it occurs less frequently.

To schedule Backup Exec to check logon accounts

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Logon Accounts**.

3 Configure any of the following options:

Check that the backup sources in jobs can be accessed with the logon accounts you selected	Select this option to test whether Backup Exec can access the backup sources in your backup jobs with the logon accounts that you selected.
Perform check every X days at X	Enter the interval for Backup Exec to test logon accounts. Backup Exec automatically tests the logon accounts at the interval that you select in this field. You can select the time at which you want Backup Exec to run the test and the number of days between tests.
Check logon accounts only at the server level	Select this option to limit the logon account test to the server level. If you select this option, Backup Exec tests only whether the logon accounts can access backup sources at the server level. Any resources that reside on the server are not tested. The logon account test takes less time if you only check accounts at the server level, however it is less thorough.

4 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Configuring Backup Exec to discover data to back up

Backup Exec’s **Discover Data to Back Up** option detects new backup content within a Windows or Active Directory domain. This option lets you configure a job that searches for new server volumes, databases, or application data.

By default, the data discovery operation runs at noon every day. It also runs each time the Backup Exec services are restarted. Backup Exec cancels the operation if it is still running after four hours. You can disable the operation or change the default settings in the global Backup Exec settings.

The **Discover Data to Back Up** option performs three main tasks:

- Discovers any top-level computers or computer contents
When the data discovery operation discovers top-level computers or computer contents, it adds them to the **Credentials** pane on the **Backup and Restore** tab. The operation updates any information about the computers or computer contents and their backup status. You can view information about backup sources on the **Credentials** pane.

- Discovers any servers that do not have an Agent for Windows installed on them
If the operation discovers any servers that do not have an Agent for Windows installed on them, Backup Exec sends you an alert. You can add the servers to the list of servers by using the **Add a Server Wizard**. After you add the servers to the list of servers, you can back them up and monitor them.
See [“Adding discovered servers to the list of servers in Backup Exec”](#) on page 715.
- Discovers and validates instances of the Agent for Windows
The data discovery operation searches for any instances of the Agent for Windows on your network. When it finds an Agent for Windows, the operation checks the version to make sure that it is up to date. If an Agent for Windows is not up to date with the most recent version, Backup Exec sends you an alert.

The data discovery operation only discovers the servers that meet the following criteria:

- Belongs to the same domain as the Backup Exec server
- Has the Windows Management Instrumentation (WMI) service enabled and running
- Allows WMI access for the same user that the Backup Exec Management Service runs under
Members of the server's "Administrators" group have this level of access.
- Has firewalls that are configured to allow WMI network traffic

To configure Backup Exec to discover data to back up

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Discover Data to Back Up**.

3 Configure any of the following options:

Discover servers that have data that has not been backed up	Select this option to allow Backup Exec to discover any data that needs to be backed up. When this option is selected, Backup Exec automatically checks your network for any data that has not been backed up.
Frequency	Lets you configure the frequency with which Backup Exec searches for any data that needs to be backed up. You can select to let Backup Exec search for any data that needs to be backed up daily, weekly, or monthly.
Interval	Lets you configure the interval at which Backup Exec searches for any data that needs to be backed up. You can select different intervals based on the frequency you selected.
Cancel data discovery if not completed within	Lets you select the number of hours after which the data discovery process is canceled if it is not finished. Canceling the data discovery process can help prevent it from affecting your system resources.

4 Click **OK**.

Adding discovered servers to the list of servers in Backup Exec

You can use the **Add Server** Wizard to install the Agent for Windows on any servers that Backup Exec discovers by the **Discover Data to Back Up** option. After you install the Agent for Windows, the server is added to the list of servers in Backup Exec.

To add discovered servers to the list of servers in Backup Exec

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Windows computers and servers**, and then click **Next**.
- 3 Select **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Click **Browse**.

- 5 Expand **Servers without an Agent for Windows installed** to see the names of the servers that the data discovery operation discovered.
- 6 Select the servers on which you want to install the Agent for Windows, and then click **OK**.
- 7 In the **Logon Account** field, select the logon account that you want to use to access each server.
- 8 Click **Next**.
- 9 Select any of the following options and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically

Upgrades the Agent for Windows on the server that you are adding to the list of servers, if necessary.

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required

Restarts the remote computer after the Agent for Windows is installed.

- 10 Click **Install**.

See [“Configuring Backup Exec to discover data to back up”](#) on page 713.

Backup networks


The backup network feature lets you direct any primary backup traffic that Backup Exec generates to a specific local network. Directing backup jobs to a specified local network isolates the backup data so that other connected networks are not affected when backup operations are performed. You also can use a backup network when you restore data. The feature is enabled on the Backup Exec server and lets you protect all the remote computers that reside on the specified local network.

When you specify a backup network and submit a job, Backup Exec verifies that the remote computer is on the same subnet as the selected interface on the Backup Exec server. If the remote computer is on the selected subnet, then the backup operation is performed.

If the remote computer is not on the selected subnet, then the job fails. However, you can set up Backup Exec to use any available network to back up remote computers.

The following diagram shows an example of a backup network configuration.

Figure 17-1 Example of backup network



In the example, the database server and mail server are connected to both the backup network and the corporate network.

When the Backup Exec server performs backup operations, the backup data uses either the backup network or the corporate network to back up the database server. If the backup data goes through the corporate network, the amount of time it takes to back up the database server increases. The amount of time increases because the network route between the two computers is longer. Users may experience network latencies when they access the mail server since there is an increase in network traffic.

In contrast, if you specify a backup network and you back up the database server, the backup data traffic is isolated to the backup network. Any users accessing the mail server are not affected. The backup network is used to perform all backup operations, unless the remote computer is not connected to the backup network.

To back up any remote computers that are not connected to the backup network, choose to use any available network route. Choosing any available network lets you back up the remote computer even though it does not reside on the backup network.

You can configure global network settings for all backup jobs on the **Network and Security** pane in the Backup Exec settings. If you want to override the global

settings for a particular backup job, you can configure network settings for individual jobs on the **Network** pane when you create backups.

See [“Changing network and security options for Backup Exec”](#) on page 718.

See [“Configuring network options for backup jobs”](#) on page 203.

See [“Using Backup Exec with firewalls”](#) on page 725.

Using IPv4 and IPv6 in Backup Exec

Backup Exec supports versions 4 and 6 of the Internet Protocol (IP), which are commonly referred to as IPv4 and IPv6. You can use IPv4 and IPv6 in backup and restore networks. Support for IPv6 is dependent upon operating system support for the protocol, as well as proper network configuration.

You can use Backup Exec in a mixed IPv4/IPv6 environment or an IPv4-only environment.

Enter an IPv4 or IPv6 address for a computer anywhere that you can enter a computer name in Backup Exec, except in the following location:

- The **Connect to Backup Exec Server** dialog box.

A Backup Exec agent that supports IPv6 can be backed up or restored using IPv6 only from a Backup Exec server that is IPv6-compliant.

Changing network and security options for Backup Exec

You can configure how Backup Exec works with your network configuration and security. The network and security options are global options that affect all Backup Exec jobs.

If the global network and security settings that you configure do not apply for a specific backup job, you can change the network settings when you create the backup job.

See [“Configuring network options for backup jobs”](#) on page 203.

To edit network and security options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Configure any of the following options:

Network interface	Select the name of the network interface card that connects the Backup Exec server to the default network that you want to use for backup jobs. The list includes all available network interfaces on the Backup Exec server.
Protocol	Select the default protocol you want to use for backup jobs. The options are as follows: ■ Use any available protocol ■ IPv4 ■ IPv6
Subnet	Select the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	Select this option to let Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network. If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.
Interface Details	Click this option to view the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.
Enable selection of user shares	Select this option to include user-defined shares in jobs. If you do not select this option, you cannot select user-defined shares when you create jobs.

Enable TCP dynamic port range

Lets Backup Exec agents use a range of ports for communication.

You enter the port range. If the first port that Backup Exec attempts to use is not available, Backup Exec attempts to use one of the other ports in the range. If none of the ports in the range is available, Backup Exec uses any available dynamic port. Default port ranges are 1025 to 65535. It is recommended that you use a range of 25 allocated ports for the remote system if you use Backup Exec with a firewall.

See [“Using Backup Exec with firewalls”](#) on page 725.

Use a custom port to receive operation requests from the Oracle server

Lets you specify the port that Backup Exec uses for communication between the Backup Exec server and the remote computer for both DBA and Backup Exec server-initiated operations. By default, Backup Exec uses port 5633.

If you change the port number on the remote Windows or Linux computer, you must also change it on the Backup Exec server. Then you must restart the Backup Exec Job Engine service on the Backup Exec server.

See [“About Oracle instance information changes”](#) on page 1243.

Use FIPS 140-2 compliant software encryption

Lets you enable software encryption that complies with FIPS 140-2 standards. If you select this option, you must use a 256-bit AES encryption key. This option is available only for Windows computers.

You must stop and restart the Backup Exec services for this change to take effect.

Manage Keys

Lets you create a new encryption key or manage existing encryption keys.

Allow only Kerberos authentication from the Remote Administration Console

Allows Backup Exec to restrict clients, such as Remote Administration Console to use Kerberos only authentication.

If you select this check box to use only Kerberos authentication, the Backup Exec Management Service must be restarted. Any connections from Remote Administration Console that fall back to use NTLM authentication are disabled.

By default, this option is not selected.

This option is disabled when a Backup Exec server is in a workgroup. Only the owner of a System Logon Account, who is logged on to the Backup Exec console has the privileges to change the option.

Note: The System Logon Account must be a member of Domain Admins group.

Secure the Backup Exec console

Lets you secure the Backup Exec console by providing the following features when you select the **Secure the Backup Exec Console** check box:

- Authentication

After you select the **Secure the Backup Exec Console** check box, the authentication setting is enabled and the next time that you launch Backup Exec, you need to enter Backup Exec login credentials to connect to the console.

If you do not enter the credentials you cannot connect to the Backup Exec console.

- Lock Console option

After you select the **Secure the Backup Exec Console** check box, this feature is enabled. You can lock the Backup Exec session that you are working on and secure the Backup Exec console from unauthorized access. Unless you unlock the Backup Exec console, you cannot perform any tasks in the Backup Exec user interface.

Note: By default, this check box is not selected.

This option is not applicable for Remote Administration Console (RAC) as you must always provide credentials to connect to the Backup Exec console.

Only the owner of a System Logon Account, who is logged on to the Backup Exec console has the privileges to change the user access settings. If you want to know the owner of the System Logon Account, click the Backup Exec button, and then select **Configuration and Settings > Select Logon Accounts > Manage Logon Accounts**. On the **Logon Account Management** dialog box, the **Owner** column displays the owner of the System Logon Account.

In case of rolling upgrade, if you have an earlier version of MMS and an updated version of CAS, and you connect to MMS from CAS, this check box is available and you can select it. However, this setting is not enabled for MMS.

See ["Locking and unlocking the Backup Exec Console"](#) on page 116.

Disk storage lockdown setting

The Ransomware Resilience feature lets you enable or disable the lockdown setting on the disk storage.

The disk storage lockdown setting protects the disk-based backup storage configured with Backup Exec. Access to disk storage is limited only to authorized processes like Backup Exec services. Only Backup Exec is allowed to write to the disk storage (Backup Exec data folders where the backup jobs are targeted). No other process can write to the disk storage. In addition, external processes are not allowed to modify backup data by injecting code into Backup Exec processes.

While the lockdown is enabled, backups and restores continue to work without a change.

If the disk storage is created on a network share hosted on a remote server, Backup Exec can only monitor the write operations originating from the media server. If the network share is accessed from any other server, which does not have Backup Exec installed, write access is allowed.

This setting is enabled by default and is the recommended setting to protect your backup data. You can disable the setting by providing the System Logon Account credentials.

See [“Disabling disk storage lockdown”](#) on page 724.

When the lockdown is in effect, the status displays **Enabled**.

The setting is disabled and the lockdown status displays **Disabled**. It is strongly recommended that you enable this setting to protect your disk-based storage. Any changes to the disk-based storage can only be done by Backup Exec.

If you disable the lockdown setting, the **Send periodic alerts if the disk storage lockdown setting is disabled** check box is selected so that you can receive periodic alerts. An alert is generated at 11 am every day until the lockdown setting is enabled again. By default, this check box is selected. You can clear the check box to stop the periodic alerts.

To enable the lockdown setting again, click **Enable**.

See [“Viewing the disk storage lockdown status”](#) on page 587.

4 Click **OK**.

If you selected the **Allow only Kerberos authentication from the Remote Administration Console** check box, a confirmation message is displayed.

5 If you want to restart the Management Service, click **Yes**.

See [“Backup networks”](#) on page 716.

Disabling disk storage lockdown

The disk storage lockdown setting is enabled by default and is the recommended setting to protect your backup data. You can disable the setting by providing the System Logon Account credentials.

To disable the disk storage lockdown

1 Click **Disable**.

The **Disable disk storage lockdown** dialog box is displayed. The disk storage lockdown setting protects the disk-based storage against unauthorised access. You must enter the Service Logon Account credentials to disable the setting.

2 Do the following:

User name	Displays the user name of the System Logon Account. You cannot edit the user name.
Password	Specify the password of the System Logon Account.
Reason to disable lockdown	Specify the reason to disable the lockdown.

3 Click **OK**.

If the setting is successfully disabled, a confirmation message is displayed. If the setting is not disabled, a message is displayed.

4 Click **OK**.

After the setting is disabled, the lockdown status displays **Disabled** on the **Network and Security** pane. It is strongly recommended that you enable this setting to protect your disk-based storage. Any changes to the disk-based storage can only be done by Backup Exec.

See [“Changing network and security options for Backup Exec”](#) on page 718.

Using Backup Exec with firewalls

In firewall environments, Backup Exec provides the following advantages:

- The number of ports that are used for backup network connections is kept to a minimum.
- Open ports on the Backup Exec server and remote systems are dynamic and offer high levels of flexibility during browsing, backup, and restore operations.
- You can set specific firewall port ranges and specify backup and restore networks within these ranges. You can use specific ranges to isolate data traffic and provide high levels of reliability.

Note: The Agent for Windows is required to perform remote backups and restores.

Firewalls affect system communication between a Backup Exec server and any remote systems that reside outside the firewall environment. You should consider special port requirements for your firewall when you configure Backup Exec.

It is recommended that you open port 10000 and make sure that it is available on the Backup Exec server and any remote systems. In addition, you must open the dynamic port ranges that Backup Exec uses for communications between the Backup Exec server and Backup Exec agents.

When a Backup Exec server connects to a remote system, it initially uses port 10000. The agent listens for connections on this predefined port. The Backup Exec server is bound to an available port, but additional connections to the agent are initiated on any available port.

When you back up data, up to two ports may be required on the computer on which the agent is installed. To support simultaneous jobs, you must configure your firewall to allow a range of ports large enough to support the number of simultaneous operations desired.

If there is a conflict, you can change the default port to an alternate port number by modifying the %systemroot%\System32\drivers\etc\services file. You can use a text editor such as Notepad to modify your NDMP entry or add an NDMP entry with a new port number. You should format the entry as follows:

```
ndmp      9999/tcp      #Network Data Management Protocol
```

Note: If you change the default port, you must change it on the Backup Exec server and all remote systems that are backed up through the firewall.

When you set up TCP dynamic port ranges, It is recommended that you use a range of 25 allocated ports for the remote computer. The number of ports that remote computers require depends on the number of devices you protect and the number of tape devices you use. You may need to increase these port ranges to maintain the highest level of performance.

Unless you specify a range, Backup Exec uses the full range of dynamic ports available. When performing remote backups through a firewall, you should select a specific range on the **Network and Security** settings dialog box.

To browse systems through a firewall

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Verify that a dynamic range of ports has been set for the Backup Exec server and the Backup Exec agent and that the firewall is configured to pass these port ranges and the 10000 port (which is used for the initial connection from the Backup Exec server to the Backup Exec agent).

Port 6101 must be open to browse Windows systems in the backup selections tree.

- 4 Click **OK**.

See [“Backup Exec ports”](#) on page 726.

See [“Backup Exec listening ports”](#) on page 728.

See [“About enabling a SQL instance behind a firewall”](#) on page 729.

Backup Exec ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a Backup Exec server and remote systems that reside outside the firewall environment.

See [“Using Backup Exec with firewalls”](#) on page 725.

The following table provides more information about which ports Backup Exec and its agents and options use:

Table 17-4 Backup Exec ports

Service or Process	Port	Port Type
Backup Exec Agent Browser (process=benetns.exe)	6101	TCP

Table 17-4 Backup Exec ports (*continued*)

Service or Process	Port	Port Type
Backup Exec Agent for Windows (process=beremote.exe)	10000	TCP
Backup Exec server (process=beserver.exe)	3527, 6106	TCP
MSSQL\$BKUPEXEC (process=sqlservr.exe)	A port number between 49152 and 65535. See Microsoft's documentation for more information on assigning the TCP/IP port numbers that are required to communicate with SQL through a firewall.	TCP UDP
Agent for Oracle on Windows or Linux Servers	Random port unless configured otherwise	
Agent for Linux and Unix	Default NDMP port, typically 10000	TCP
Backup Exec deduplication engine (process=spoold.exe)	10082	TCP
Backup Exec deduplication manager (process=spad.exe)	10102	TCP
Backup Exec Management Service (process=BackupExecManagementService.exe)	50104	TCP
Kerberos	88	UDP
NETBIOS	135	TCP, UDP
NETBIOS Name Service	137	UDP
NETBIOS Datagram Service	138	UDP
NETBIOS Session Service	139	TCP
NETBIOS	445	TCP

Table 17-4 Backup Exec ports (*continued*)

Service or Process	Port	Port Type
DCOM/RPC	3106	TCP
Agent for Windows	6103	TCP
Push Install - Check for conflicts in message queue for CAS, which is part of beserver.exe	Random port number between 1030 and 1039	TCP
Push Install	441	TCP
SMTP email notification	25 587 465 Outbound from Backup Exec server	TCP
SNMP	162 outbound from Backup Exec server	TCP

Backup Exec listening ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a Backup Exec server and remote systems that reside outside the firewall environment.

See [“Using Backup Exec with firewalls”](#) on page 725.

When Backup Exec is not running operations, it listens to ports for incoming communication from other services and agents. Backup Exec initially communicates with the agent using a static listening port to begin an operation. The agent and the Backup Exec server then use dynamic ports to pass data back and forth.

Backup Exec uses the following listening ports:

Table 17-5 Backup Exec listening ports

Service	Port	Port Type
Backup Exec Agent Browser (benetns.exe)	6101	TCP
Backup Exec Agent for Windows (beremote.exe)	10000	TCP

Table 17-5 Backup Exec listening ports (continued)

Service	Port	Port Type
Backup Exec server (beserver.exe)	3527, 6106	TCP
Backup Exec Management Service (process= BackupExecManagementService.exe)	50104	TCP
MSSQL\$BKUPEXEC (sqlsevr.exe)	1125	TCP
	1434	UDP
Agent for Linux and Unix (RALUS)	10000	TCP
DBA-initiated backups for Oracle	5633	TCP

About enabling a SQL instance behind a firewall

If you want to connect to a SQL instance behind a firewall, you must enable the SQL instance for communication. To enable the SQL instance for communication, you must make the SQL port static and configure the Windows Firewall.

The Backup Exec SQL instance is configured to use a dynamic port by default. Each time SQL Server is started, the port number can change.

You also must configure the Windows Firewall to allow connections to the SQL instance. There may be multiple ways to configure the Windows Firewall based on your system configuration. You can add sqlsvr.exe and sqlbrowser.exe to the Windows Firewall Exceptions list or you can open a port in the Windows Firewall for TCP access. Refer to the Microsoft Knowledge Base for more information or to determine which configuration is best for your network.

See [“Installing a managed Backup Exec server across a firewall ”](#) on page 1333.

See [“Using Backup Exec with firewalls”](#) on page 725.

Using encryption with Backup Exec

Backup Exec provides you with the ability to encrypt data. When you encrypt data, you protect it from unauthorized access. Anyone that tries to access the data has to have an encryption key that you create. Backup Exec provides software encryption, but it also supports some devices that provide hardware encryption with the T10 standard. Backup Exec configures encryption when you specify which storage devices that you want to use for a backup job.

Backup Exec supports two security levels of encryption: 128-bit Advanced Encryption Standard (AES) and 256-bit AES. Within 256-bit AES encryption level, Backup Exec provides two methods of key derivation for the pass phrase that you entered. One is the existing AES 256-bit that uses SHA-2 algorithm (earlier referred to as 256-bit AES) and the second makes use of PBKDF2 that is the enhanced password-based Key Derivation Function algorithm.

The 256-bit AES encryption provides a stronger level of security because the key is longer for 256-bit AES than for 128-bit AES. However, 128-bit AES encryption enables backup jobs to process more quickly. Hardware encryption using the T10 standard requires 256-bit AES. With PBKDF2, in addition to pass phrase, Backup Exec uses randomly generated Salt, which makes the encryption key more secure.

When you run a duplicate backup job, any backup sets that are already encrypted, will remain encrypted, regardless of the encryption option that you select. However, you can encrypt any unencrypted backup sets.

For information about the best practices of Backup Exec software encryption, refer to *Backup Exec Best Practices*.

This topic includes the following information:

[Software encryption](#)

[Hardware encryption](#)

[Encryption keys](#)

[Restricted keys and common keys](#)

[Pass phrases](#)

Software encryption

When you install Backup Exec, the installation program installs encryption software on the Backup Exec server and on any remote computers that use a Backup Exec agent. Backup Exec can encrypt data at a computer that uses a Backup Exec agent, and then transfer the encrypted data to the Backup Exec server. Backup Exec then writes the encrypted data on a set-by-set basis to tape or to disk storage.

Backup Exec encrypts the following types of data:

- User data, such as files and Microsoft Exchange databases.
- Metadata, such as file names, attributes, and operating system information.
- On-tape catalog file and directory information.

Backup Exec does not encrypt Backup Exec metadata or on-disk catalog file and directory information.

You can use software compression with encryption for a backup job. First Backup Exec compresses the files, and then encrypts them. However, backup jobs take longer to complete when you use both encryption compression and software compression.

It is recommended that you avoid using hardware compression with software encryption. Hardware compression is performed after encryption. Data becomes randomized during the encryption process. Compression does not work effectively on data that is randomized.

Hardware encryption

Backup Exec supports hardware encryption for any storage devices that use the T10 encryption standard. When you use hardware encryption, the data is transmitted from the host computer to the storage device and then encrypted on the device. Backup Exec manages the encryption keys that are used to access the encrypted data.

Backup Exec only supports approved devices for T10 encryption.

Note: Hardware encryption that uses the T10 standard requires 256-bit AES. Backup Exec does not let you enable hardware encryption for a job unless it uses at least a 16-character pass phrase.

Encryption keys

You must create encryption keys to use encryption in Backup Exec. When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user's security identifier. The person who creates the key becomes the owner of the key.

If you use encryption for synthetic backups, all of the associated backups must use the same encryption key. Do not change the encryption key after the baseline is created. The encryption key that you select for the baseline backup is automatically applied to all associated backups.

When you select encrypted data for restore, Backup Exec verifies that encryption keys for the data are available in the database. If any of the keys are not available, Backup Exec prompts you to recreate the missing keys. If you delete the key after you schedule the job to run, the job fails.

If Backup Exec cannot locate an encryption key while a catalog job is running, Backup Exec sends an alert. You can then recreate the missing encryption key if you know the pass phrase. If the Backup Exec alert contains Salt information, you must provide the same salt to recreate the missing encryption key.

Simplified Disaster Recovery supports the recovery of computers with previously encrypted backup sets. If you have Simplified Disaster Recovery backups that are encrypted during backup, the **Recover This Computer** wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the recovery.

See [“Encryption key management”](#) on page 733.

Restricted keys and common keys

Backup Exec has the following types of encryption keys:

Table 17-6 Types of encryption keys

Key type	Description
Common	Anyone can use the key to encrypt data during a backup job and to restore encrypted data.
Restricted	Anyone can use the key to encrypt data during a backup job, but users other than the key owner must know the pass phrase. If a user other than the key owner tries to restore the encrypted data, Backup Exec prompts the user for the pass phrase. If you cannot supply the correct pass phrase for the key, you cannot restore the data.

Pass phrases

Encryption keys require a pass phrase, which is similar to a password. Pass phrases are usually longer than passwords and are comprised of several words or groups of text. A good pass phrase is between 8 and 128 characters. The minimum number of characters for 128-bit AES encryption is eight. The minimum number of characters for 256-bit AES encryption (SHA-2) and 256-bit AES encryption (PBKDF2) is 16. It is recommended that you use more than the minimum number of characters.

Note: Hardware encryption that uses the T10 standard requires 256-bit AES. Backup Exec does not let you enable hardware encryption for a job unless it uses at least a 16-character pass phrase.

Also, a good pass phrase contains a combination of upper and lower case letters, numbers, and special characters. You should avoid using literary quotations in pass phrases.

For 256-bit AES PBKDF2, the pass phrase must contain at least one upper case, one lower case, one number, and one special character.

A pass phrase can include only printable ASCII characters, which are characters 32 through 126. ASCII character 32 is the space character, which is entered using the space bar on the keyboard. ASCII characters 33 through 126 include the following:

```
!"#$%&'()*+,-./0123456789;<=>?@ABCDEFGHIJKLMN OPQRSTUVWXYZ  
[\]^_`abcdefghijklmnopqrstuvwxyz{|}~
```

See [“Encryption key management”](#) on page 733.

Encryption key management

When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user’s security identifier. The person who creates the key becomes the owner of the key.

See [“Using encryption with Backup Exec”](#) on page 729.

Backup Exec stores the keys in the Backup Exec database. However, Backup Exec does not store the pass phrases for the keys. The owner of each key is responsible for remembering the pass phrase for the key.

To protect your keys, it is recommended that you do the following:

- Maintain a written log of the pass phrases. Keep the log in a safe place in a separate physical location from the encrypted backup sets.
- Back up the Backup Exec database. The database keeps a record of the keys.

Caution: If you do not have a backup of the Backup Exec database and do not remember your pass phrases, you cannot restore data from the encrypted media.

A key that is created on a Backup Exec server is specific to that Backup Exec server. You cannot move keys between Backup Exec servers. However, you can create new keys on a different Backup Exec server by using existing pass phrases. A pass phrase always generates the same key. In addition, if you delete a key accidentally, you can recreate it by using the pass phrase.

When you perform the catalog operation, Backup Exec sends an alert with the Salt. Use this Salt and the same or original pass phrase to recreate the same or missing encryption key again.

See [“Creating encryption keys”](#) on page 734.

If a Backup Exec database becomes corrupted on a Backup Exec server and is replaced by a new database, you must manually recreate all of the encryption keys that were stored on the original database.

If you move a database from one Backup Exec server to another Backup Exec server, the encryption keys remain intact as long as the new Backup Exec server meets the following criteria:

- Has the same user accounts as the original Backup Exec server.
- Is in the same domain as the original Backup Exec server.

See [“Replacing an encryption key”](#) on page 735.

See [“Deleting encryption keys”](#) on page 738.

Creating encryption keys

When you create an encryption key, you select the type of encryption to use.

To create an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Click **New**.
- 5 In the **Key name** field, type a unique name for this key. The name can include up to 256 characters.
- 6 In the **Encryption type** field, select the encryption type to use for this key. Your choices are **128-bit AES**, **256-bit AES (SHA-2)**, or **256-bit AES (PBKDF2)**. 256-bit AES (SHA-2) was earlier known as 256-bit AES.

The default type is 256-bit AES PBKDF2. The 256-bit AES encryption provides a stronger level of security than 128-bit AES encryption. However, backup jobs may process more slowly with 256-bit AES encryption than with 128-bit AES encryption. Hardware encryption that uses the T10 standard requires 256-bit AES.

- 7 In the **Pass phrase** field, type a pass phrase for this key. You can use only printable ASCII characters.

For 128-bit AES encryption, the pass phrase must be at least eight characters. For 256-bit AES encryption, the pass phrase must be at least 16 characters.

It is recommended that you use more than the minimum number of characters. For 256-bit AES PBKDF2, the pass phrase must have at least one upper case, one lower case, one number, and one special character.

Warning: If an encryption key that is used in a backup is no longer available, you must provide the pass phrase during restore. Without the pass phrase, the data cannot be accessed.

- 8 In the **Confirm pass phrase** field, type the pass phrase again to confirm it.
- 9 (Optional) Select the check box and then enter the Salt to recreate the same or missing encryption key again.

If Backup Exec cannot find the encryption key in the database during restore or cataloging, then you need to create the same key again. If the key is created using 256-bit AES PBKDF2, you must enter the Salt. Ensure that you only enter the Salt given by Backup Exec. Salt information would be displayed in a Backup Exec alert during catalog operation.

- 10 In the **Encryption key type** group box, select whether you want to create a common or restricted encryption key.

If a key is common, any user of this installation of Backup Exec can use the key to back up and restore data. If a key is restricted, anyone can use the key to back up data. But only the key owner or a user who knows the pass phrase can use the restricted key to restore the encrypted data.

- 11 Click **OK**.

See [“Encryption key management”](#) on page 733.

See [“Using encryption with Backup Exec”](#) on page 729.

Replacing an encryption key

You can replace one encryption key with another for all backup jobs and duplicate backup set jobs.

Note: You cannot replace an encryption key if it is used in a restore job.

To replace an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Select the key that you want to replace.
- 5 Click **Replace**.
- 6 In the **Select an encryption key to replace <key name>** field, do one of the following:

To use an existing key Select the key from the list.

To create a new key

Complete the following steps:

- Click **New**.
- In the **Key name** field, type a unique name for the key.
The name can include up to 256 characters.
- In the **Encryption type** field, select the type of encryption to use for this key.
You can select **128-bit AES**, **256-bit AES (SHA-2)**, or **256-bit AES (PBKDF2)**. 256-bit AES (SHA-2) was earlier known as 256-bit AES. The default type is 256-bit AES PBKDF2.
The 256-bit AES encryption provides a stronger level of security than 128-bit AES encryption. However, backup jobs may process more slowly with 256-bit AES encryption than with 128-bit AES encryption.
Hardware encryption that uses the T10 standard requires 256-bit AES.
- In the **Pass phrase** field, type a pass phrase for this key.
For 128-bit AES encryption, the pass phrase must be at least eight characters. For 256-bit AES encryption, the pass phrase must be at least 16 characters. It is recommended that you use more than the minimum number of characters. For 256-bit AES PBKDF2, the pass phrase must have at least one upper case, one lower case, one number, and one special character.
You can use only printable ASCII characters.
Warning: If an encryption key that is used in a backup is no longer available, you must provide the pass phrase during restore. Without the pass phrase, the data cannot be accessed.
- In the **Confirm pass phrase** field, type the pass phrase again to confirm it.
- (Optional) Select the check box and then enter the Salt to recreate the same or missing encryption key again.
If Backup Exec cannot find the encryption key in the database during restore or cataloging, then you need to create the same key again. If the key is created using 256-bit AES PBKDF2, you must enter the Salt. Ensure that you only enter the Salt given by Backup Exec. Salt information would be displayed in a Backup Exec alert during catalog operation.
- In the **Encryption key type** group box, select whether you want to create a common or restricted encryption key.
- Click **OK**.

7 Click **OK**.

See [“Encryption key management”](#) on page 733.

See [“Using encryption with Backup Exec”](#) on page 729.

Deleting encryption keys

You should be cautious when you delete encryption keys. When you delete an encryption key, you cannot restore the backup sets that you encrypted with that key unless you create a new key that uses the same encryption key and pass phrase as the original key.

You can delete encryption keys in the following situations:

- The encrypted data on the tape has expired or the tape is retired.
- The encryption key is not the default key.
- The encryption key is not being used in a job. If the key is being used, you must select a new key for the job.

If you delete an encryption key that is being used in a scheduled restore job, you cannot replace the key. Therefore, any scheduled restore job in which you delete an encryption key fails.

To delete an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Select the key that you want to delete.
- 5 Click **Delete**.
- 6 Click **Yes**.
- 7 If the key is used in a job, do the following:
 - In the **Select an encryption key to replace "key name"** box, select the new key for the job or click **New** to create a new key.
 - Click **OK**.

See [“Encryption key management”](#) on page 733.

See [“Using encryption with Backup Exec”](#) on page 729.

Encryption keys and Salt

When you create or delete a new data encryption key, information about all keys is exported to a csv in the Backup Exec Data (BEData) folder. The csv file is named `EncryptionKeys-<MediaServerMachineName>.csv`.

For a PBKDF2 key, the combination of key name, Passphrase, and its associated Salt is required. You can use this option to view the Salt for a PBKDF2 key. Backup Exec requires the Passphrase and the Salt combination if the key is not available on the media server.

The `EncryptionKeys-<MediaServerMachineName>.csv` file contains the following information:

- **ProductVersion:** Backup Exec version installed
- **MachineName:** Media server name
- **Key:** Contains the following information:
 - Name of the key
 - Type of key: 128-bit AES, 256-bit AES (SHA-2), or 256-bit AES (PBKDF2)
 - Salt Length of the key: Number of characters
 - Salt value
 - Restricted key or common encryption key

It is recommended that you keep a backup of the csv file, which may be required during a disaster recovery scenario.

Granular Recovery Technology

You can use Granular Recovery Technology (GRT) to restore certain individual items from backup sets. For example, you can use the Agent for Microsoft Exchange Server to restore an email from a backup without having to restore the entire mailbox. Or, you can use the Agent for Microsoft SharePoint to restore a list without restoring the entire site.

To restore individual items, the Granular Recovery Technology feature must be enabled when you create a backup job.

GRT is enabled by default for backups for the following agents:

- Agent for Microsoft Active Directory
- Agent for Microsoft Exchange Server
- Agent for Microsoft SharePoint

■ Agent for VMware and Hyper-V

You can restore either full backup sets or individual items from GRT-enabled backups.

By default, the Agent for VMware and Hyper-V uses Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SharePoint, and Active Directory application data that resides on virtual machines.

When you back up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. GRT-enabled backup jobs require more time to catalog because of the amount of granular information that they contain. For GRT-enabled backup jobs, by default the catalog operation is delayed and runs as a separate operation to have less effect on your backup window. Because the catalog operation runs separately from the backup job, it does not prevent another scheduled backup job from starting on time.

When you enable GRT for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware backups, the full catalog operation runs immediately after the backup job by default.

For Exchange and SharePoint agent-based backups, the full catalog operation runs immediately after all full backups. It runs once every 24-hours for all incremental backups and differential backups, even if you schedule more than one GRT-enabled job to run in the 24 hour period.

For Hyper-V and VMware backups, the full catalog operation runs immediately after all full, incremental, and differential backups by default.

You can configure the full catalog operation to run on a schedule if you do not want it to run immediately after the backup job. You can also run an Instant GRT operation that runs as part of the backup job.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

For information about the best practices to use Backup Exec and Granular Recovery Technology (GRT), refer to *Backup Exec Best Practices*.

The following table lists the individual items you can restore for each agent.

Table 17-7 Individual items that can be recovered for each agent

Agent	Individual items
Agent for Microsoft Active Directory	You can restore the following individual items: ■ Active Directory objects and attributes ■ Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes
Agent for Microsoft Exchange Server	You can restore the following individual items: ■ Mailboxes ■ Mail messages and their attachments ■ Public folders ■ Calendar items ■ Contacts ■ Notes ■ Tasks
Agent for Microsoft SharePoint	The following are examples of the individual items that can be restored: ■ Site collections ■ Sites or subsites ■ Document or picture libraries ■ Lists ■ Individual list items ■ Documents, pictures, or other files that are stored in libraries
Agent for VMware and Hyper-V	You can restore drives, folders, and files from virtual machines that run a Windows operating system. You can also enable the granular recovery of Microsoft Exchange, SQL, SharePoint, and Active Directory application data that resides on virtual machines. See “Using Granular Recovery Technology (GRT) with the Agent for VMware” on page 1028.

When you run a GRT-enabled backup job, Backup Exec creates media with an IMG prefix (for example, IMG00001). IMG media is a specific media type that Backup Exec creates only for GRT-enabled backup operations. When you run a GRT-enabled backup job, the IMG media stores the backup data.

You should consider which device you use for GRT-enabled backups before you begin. You should also consider any special requirements for the type of data you back up.

Recommended devices for backups that use Granular Recovery Technology

It is recommended that you select a disk storage device for any backups that are enabled for Granular Recovery Technology (GRT). The disk storage device should be on a volume that does not have file size limitations. An NTFS drive is an example of a volume without file size limitations. Some examples of volumes that have file size limitations include FAT and FAT32 volumes.

Note: Backup Exec does not store the granular backup sets on disk in encrypted form when you enable encryption for the GRT-enabled backup jobs that are sent to disk, deduplication, and disk cartridge devices. Only the backup sets for the backup sources that do not support GRT are stored in encrypted form. All the backup sets for the backup jobs that are sent to cloud, OpenStorage, and tape devices are stored in encrypted form.

If you must use a disk storage device on a volume with file size limitations, Backup Exec requires a staging location. Backup Exec temporarily stores a small amount of metadata in the staging location during the backup job. It deletes the data from the staging location when the backup is finished. The staging location is not necessary, however, if you use a disk storage device on a volume without file size limitations as the destination.

The staging location's default path is C:\temp.

The volume that is used for a staging location for backup jobs should meet the following requirements:

- It is local to the Backup Exec server
- It does not have any file size limitations

Additionally, to avoid disk space problems, it should meet these requirements:

- It should not be a system volume
- It should have at least 1 GB of available space

Backup Exec also uses a staging location to restore GRT-enabled data from a tape or from a disk storage device on volumes with file size limitations. The staging location must be on a volume that does not have file size limitations and is local to the Backup Exec server. The staging location is not necessary if you restore GRT-enabled data from disk storage on a volume without file size limitations, such as NTFS.

Backup Exec uses the staging area differently for the following types of restores:

Table 17-8 Staging processes

Location of data to be restored	Staging process
Tape	Backup Exec copies the entire backup set or sets to the staging area. The staging area must have enough disk space for the entire backup set or sets from which you want to restore an individual item. Before you use a tape device for a GRT-enabled backup, ensure that sufficient disk space is available to perform a restore. Backup Exec deletes the data from the staging area when the restore job is complete.
Cloud	Backup Exec uses a staging location to restore GRT-enabled data from the Cloud Storage Device. Backup Exec copies the entire backup set or sets to the staging area. The staging area must have enough disk space for the entire backup set or sets from which you want to restore an individual item. Before you use a Cloud Storage Device for a GRT-enabled backup, ensure that sufficient disk space is available to perform a restore. Backup Exec deletes the data from the staging area when the restore job is complete.
Disk storage device that is on a volume with file size limitations (such as FAT or FAT32)	Backup Exec must copy a small amount of metadata that is associated with the backup set to the staging area to complete the restore. Backup Exec deletes the data from the staging area when the restore job is complete.

The staging location's default path is C:\temp. You can change the default backup and restore staging locations in the Backup Exec settings.

Requirements for jobs that use Granular Recovery Technology

Keep in mind the following requirements when you use Granular Recovery Technology (GRT) with the agents listed:

Table 17-9 Granular Recovery Technology requirements

Agent	Restrictions
Agent for Microsoft Active Directory	You must use a Backup Exec server that is running Windows Server 2012 R2 to back up an Active Directory server that is running Windows Server 2012 R2. In a CAS environment, Backup Exec runs the backup job on the central administration server if it is running Windows Server 2012 R2 and the storage is configured locally. Otherwise, Backup Exec attempts to find a managed Backup Exec server that is running Windows Server 2012 R2 to run the job. If it cannot find one, the job completes with a status of success with exceptions and the backup sets are not GRT-enabled. If the Active Directory server is a virtual machine, Backup Exec may not be able to detect that the server is running Windows Server 2012 R2. The resulting backup may not be GRT-enabled.
Agent for Microsoft Exchange Server	Backup Exec must have access to a uniquely named mailbox within the Exchange organization for backup and restore of the Information Store. See “Requirements for accessing Exchange mailboxes ” on page 1175. Backup Exec uses a disk storage device that does not have file size limitations and is local to the Backup Exec server as the default staging location for GRT-enabled backups of Exchange. Backup Exec may use a staging location other than the one that is configured in default backup options because a disk that uses the same disk geometry as the database logs is required to perform GRT operations. However, the location does not affect the performance or the amount of disk space that is required for the operations. You can create a vhd file or vhdx file to use as the staging area instead of a physical volume, if you want to minimize the effect on your resources. Configure the vhd or the vhdx file as a large-sector volume and make sure that it is at least 1GB in size. Then mount the vhd file or vhdx file to a directory. Backup Exec automatically finds the virtual volume when it needs a staging area to back up any data that resides on a large-sector drive. Recommended devices for backups that use Granular Recovery Technology

Table 17-9 Granular Recovery Technology requirements (continued)

Agent	Restrictions
Agent for Microsoft SharePoint	You must have a current version of the Agent for Windows installed on all of the servers that participate in the SharePoint farm.
Agent for VMware and Hyper-V	You can recover only individual items to virtual machines that run a Windows operating system. By default, the Agent for VMware and Hyper-V uses Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SQL, SharePoint, and Active Directory application data that resides on virtual machines. If you back up a virtual machine to tape, the Use storage-based catalogs option must be enabled in the Catalog settings to ensure that you can recover individual items from the backup sets. See “Configuring default options for catalogs” on page 255.

See [“Setting default Granular Recovery Technology \(GRT\) options”](#) on page 745.

Setting default Granular Recovery Technology (GRT) options

Backup Exec's Granular Recovery Technology (GRT) feature lets you restore certain individual items from backup sets. For example, you can use the Agent for Microsoft Exchange Server to restore an email from a backup without having to restore the entire mailbox. Or, you can use the Agent for Microsoft SharePoint to restore a list without restoring the entire site.

See [“Granular Recovery Technology”](#) on page 739.

You can configure default settings for Granular Recovery Technology. Backup Exec applies the default settings to any backup jobs that you configure to use GRT.

To set default Granular Recovery Technology options

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Granular Recovery Technology**.

3 Configure any of the following options:

If Granular Recovery Technology (GRT) is enabled for backups, enter the path to an NTFS volume of the local Backup Exec server where Backup Exec can stage temporary data

Lets you designate a location where Backup Exec can stage temporary data during GRT-enabled backup jobs.

Ensure that the location is an NTFS volume and that it is not a system volume. If the default path of C:\TEMP does not meet these requirements, type a different path on the Backup Exec server where Backup Exec can stage temporary data.

Backup Exec deletes the data when the backup job is completed.

At least 1 GB of disk space is required.

Enter the path to an NTFS volume of the local Backup Exec server where Backup Exec can store temporary data (Microsoft Hyper-V, Microsoft Exchange, Microsoft SharePoint, Microsoft Active Directory, and VMware)

Lets you designate a location where Backup Exec can stage temporary data during GRT restore jobs.

This option is applicable only when you restore individual items under the following conditions:

- The backup of Microsoft Hyper-V, Microsoft Exchange, Microsoft SharePoint, Microsoft Active Directory, or VMware Virtual Infrastructure was enabled for Backup Exec GRT.
- The backup is on a tape.
- The backup is on disk storage on a volume that has size limitations. FAT and FAT32 are examples of types of volumes that have file size limitations.

Type the path to a folder on an NTFS volume on this Backup Exec server. Restore data and metadata for this job are stored here temporarily before the individual items are restored. The staged data is automatically deleted when the restore job is completed.

It is recommended that you avoid using system volumes for temporary staging locations.

4 Click OK.

DBA-initiated job templates

When you create a DBA-initiated backup operation, you can specify the default job template in Backup Exec. You can also specify a new job template that you create in Backup Exec. The job template contains the settings that Backup Exec applies to DBA-initiated jobs.

Make sure that the name of the job template that you want to use is also configured in the instance information on the Windows computer.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1246.

Note the following about DBA-initiated jobs:

- DBA-initiated jobs fail when the related job template is deleted. To stop DBA-initiated jobs from running, delete the related DBA-initiated job template.
- All DBA-initiated backup and restore jobs are deleted after the jobs have completed.
- You cannot set minimum device requirements for DBA-initiated jobs.

See [“Creating DBA-initiated job templates”](#) on page 747.

See [“Editing DBA-initiated job templates”](#) on page 748.

See [“Deleting DBA-initiated job templates”](#) on page 748.

Creating DBA-initiated job templates

You can create a new job template that Backup Exec applies to DBA-initiated jobs.

To create DBA-initiated job templates

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Click **New**.
- 4 In the left pane, select the type of options that you want to set. You determine the options that need to be set based on the needs of your environment.
- 5 Click **OK**.

See [“DBA-initiated job templates”](#) on page 746.

See [“Storage options for DBA-initiated jobs”](#) on page 749.

See [“General options for DBA-initiated jobs”](#) on page 754.

See [“Network options for DBA-initiated jobs”](#) on page 754.

See [“Migrator for Enterprise Vault options”](#) on page 1301.

See [“Notification options for jobs”](#) on page 322.

See [“Duplicate job settings for DBA-initiated jobs”](#) on page 755.

Editing DBA-initiated job templates

You can edit the job template settings that Backup Exec applies to DBA-initiated jobs.

To edit DBA-initiated job templates

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Select the job template that you want to edit.
- 4 Click **Edit**.
- 5 In the left pane, select the type of options that you want to edit. You determine the options that need to be set based on the needs of your environment.
- 6 Click **OK**.

See [“DBA-initiated job templates”](#) on page 746.

See [“Storage options for DBA-initiated jobs”](#) on page 749.

See [“General options for DBA-initiated jobs”](#) on page 754.

See [“Network options for DBA-initiated jobs”](#) on page 754.

See [“Migrator for Enterprise Vault options”](#) on page 1301.

See [“Notification options for jobs”](#) on page 322.

See [“Duplicate job settings for DBA-initiated jobs”](#) on page 755.

Deleting DBA-initiated job templates

You can delete the templates that Backup Exec applies to DBA-initiated jobs if you no longer use them.

To delete a job template for DBA-initiated jobs

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Select the job template that you want to delete.

- 4 Click **Delete**.
- 5 Click **Yes**.
- See [“DBA-initiated job templates”](#) on page 746.

Storage options for DBA-initiated jobs

You can configure storage settings for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 746.

Table 17-10 Storage options for DBA-initiated jobs

Item	Description
Storage	Specifies the storage device to which you want to send backup data for the DBA-initiated jobs. See “Creating storage device pools” on page 527. See “About the Remote Media Agent for Linux ” on page 1456. See “Features and types of disk-based storage and network-based storage” on page 330.
Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported	Enables a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device, and to perform client-side deduplication if the device supports it. The Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If client-side deduplication cannot be performed, then either Backup Exec server deduplication or Appliance deduplication is performed. This option appears if the Deduplication feature is installed and an OpenStora~~e~~ device or a deduplication disk storage device is selected in the Storage field. See “How to use client-side deduplication” on page 995.

Table 17-10 Storage options for DBA-initiated jobs *(continued)*

Item	Description
Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported	Enables a remote computer to send data through the Backup Exec server to an OpenStorage device or a deduplication disk storage device, and to perform Backup Exec server-side deduplication if it is supported. If the Backup Exec server does not support deduplication, the data is deduplicated on an intelligent disk device, such as PureDisk or a device from a third-party vendor. This option appears if the Deduplication feature is installed and an OpenStorage device or a deduplication disk storage device is selected in the Storage field. See “About the Deduplication feature” on page 969.
Keep for	Designates the amount of time for which you want to keep the backup sets or job history from the DBA-initiated jobs.
Media set	Indicates the media set to use for the DBA-initiated jobs. The media set specifies the overwrite protection period and the append period for the backup data on the media. If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu. This option is available only if you selected a tape device in the Storage field. See “Default media sets” on page 491.

Table 17-10 Storage options for DBA-initiated jobs (continued)

Item	Description
Overwrite media	Indicates that the backup job is placed on an overwritable media. Ensure that appropriate media is in the storage device that you select. Appropriate media for an overwrite job includes the following: ■ Scratch media ■ Media for which the overwrite protection period has expired Allocated or imported media may also be overwritten depending on the media overwrite protection level that is set. Depending on your configuration, overwritable media is selected from scratch media or recyclable media. If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media. This option is available only if you selected a tape device in the Storage field. See “Managing tapes” on page 491. See “Media overwrite protection levels for tape media” on page 505. See “How Backup Exec searches for overwritable media in tape drives ” on page 506.
Append to media, overwrite if no appendable media is available	Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set. If an append job fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media. This option is available only if you selected a tape device in the Storage field.

Table 17-10 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Append to media, terminate job if no appendable media is available	Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job. This option is available only if you selected a tape device in the Storage field.
Eject media after job completes	Ejects the media from the drive or slot when the operation completes. You can also schedule a job to eject media. This option is available only if you selected a tape device in the Storage field. See “Ejecting media from a disk cartridge or tape drive” on page 567.
Retension media before backup	Runs the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.
Use Write once, read many (WORM) media	Specifies the use of WORM (write once, read many) media as the default for DBA-initiated jobs. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent. See “How WORM media is used in Backup Exec” on page 512.

Table 17-10 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Compression	Provides the following compression options: ■ None Copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space. Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive. ■ Software Uses STAC software data compression, which compresses the data before it is sent to the storage device. ■ Hardware (if available, otherwise none) Uses hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed. ■ Hardware (if available, otherwise software) Uses hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.
Encryption type	Specifies the type of encryption that you want to use, if any. See “Using encryption with Backup Exec” on page 729.
Encryption key	Specifies the encryption key that you want to use, if you selected to use encryption. See “Using encryption with Backup Exec” on page 729.

Table 17-10 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Manage Keys	Lets you manage your encryption keys. You can delete or replace existing encryption keys. You can also create a new encryption key. This option is available only if you select an encryption type. See “Encryption key management” on page 733.

General options for DBA-initiated jobs

You can configure general options for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 746.

Table 17-11 General options for DBA-initiated jobs

Item	Description
Job name	Specifies the name for this backup template. You can accept the default name that appears or enter a name. The name must be unique.
Backup set description	Describes the information in the backup set for future reference.
Verify after backup completes	Performs a verify operation automatically to make sure that the media can be read after the backup has been completed. Verifying all backups is recommended.

Network options for DBA-initiated jobs

You can configure network options for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 746.

Note: Some of these options may not display in a CAS environment.

Table 17-12 Network options for DBA-initiated jobs

Item	Description
Network interface	Specifies the name of the network interface card that connects the Backup Exec server to the network that you want to use for this backup job. The list includes all available network interfaces on the Backup Exec server.
Protocol	Specifies the protocol you want to use for this backup job. The options are as follows: ■ Use any available protocol ■ Use IPv4 ■ Use IPv6
Subnet	Displays the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	Lets Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network. If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.
Interface Details	Displays the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.
Allow managed Backup Exec server to use any network interface to access Backup Exec agents	Lets a job use any network interface to access Backup Exec agents if the selected network interface is unavailable. Enabling this option lets the managed Backup Exec server use an alternate network interface to run any important backup jobs that would otherwise fail. This option is available only if the Central Admin Server feature (CAS) is installed. See “About the Central Admin Server feature” on page 1321.

Duplicate job settings for DBA-initiated jobs

You can configure duplicate job template settings for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 746.

Table 17-13 Duplicate job settings for DBA-initiated jobs

Item	Description
Enable settings to duplicate backup sets for this job	Enables the settings for a duplicate backup set template.
Storage	Specifies the storage device to which you want to send backup data for the duplicate DBA-initiated job.
Keep for	Designates the amount of time for which you want to keep the backup sets or job history from the duplicate DBA-initiated job.
Media set	Indicates the media set to use for the duplicate DBA-initiated job. The media set specifies the overwrite protection period and the append period for the backup data on the media. If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu. See “Default media sets” on page 491.
Overwrite media	Indicates that the backup job is placed on an overwritable media. Ensure that appropriate media is in the storage device that you select. Appropriate media for an overwrite job include the following: ■ Scratch media ■ Media for which the overwrite protection period has expired Depending on your configuration, overwritable media is selected from scratch media or recyclable media. If the media in the storage device is not overwritable, an alert prompts you to insert overwritable media. See “Managing tapes” on page 491. See “Media overwrite protection levels for tape media” on page 505. See “How Backup Exec searches for overwritable media in tape drives ” on page 506.

Table 17-13 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Append to media, overwrite if no appendable media is available	Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set. If an append job fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.
Append to media, terminate job if no appendable media is available	Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job.
Eject media after job completes	Ejects the media from the drive or slot when the operation completes. You can also schedule a job to eject media. See “Ejecting media from a disk cartridge or tape drive” on page 567.
Retension media before backup	Runs the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.
Use Write once, read many (WORM) media	Specifies the use of WORM (write once, read many) media as the default for DBA-initiated jobs. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent. See “How WORM media is used in Backup Exec” on page 512.
Enable DirectCopy to tape	Enables Backup Exec to coordinate the movement of data from virtual storage directly to a physical storage device. The Backup Exec server records information about the data in the catalog. Therefore, you can restore data from either the virtual storage or the physical storage. See “Copying data from a virtual tape library to a physical tape device using DirectCopy to tape” on page 229.

Table 17-13 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Compression	Provides the following compression options: ■ None Copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space. Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive. ■ Software Uses STAC software data compression, which compresses the data before it is sent to the storage device. ■ Hardware (if available, otherwise none) Uses hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed. ■ Hardware (if available, otherwise software) Uses hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.
Encryption type	Specifies the encryption key that you want to use, if any. See “Using encryption with Backup Exec” on page 729.
Encryption key	Specifies the encryption key that you want to use, if you selected to use encryption. See “Using encryption with Backup Exec” on page 729.

Table 17-13 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Manage Keys	Lets you manage your encryption keys. You can delete or replace existing encryption keys. You can also create a new encryption key. This option is available only if you select an encryption type. See “Encryption key management” on page 733.
Preferred source device	Specifies the preferred source device that you want to use as the storage for the duplicate job.
Verify after backup completes	Performs a verify operation automatically to make sure that the data can be read after the backup has been completed. Verifying all backups is recommended.

Backup Exec logon accounts

A Backup Exec logon account stores the credentials of a user account that you use to access a computer. Backup Exec logon accounts enable Backup Exec to manage user names and passwords and can be used to browse computers or process jobs. Using Backup Exec logon accounts enables you to apply credential changes to the jobs that use them.

Backup Exec logon accounts are used to browse local and remote computers. Whenever the Backup Exec logon credentials are passed between the Backup Exec server and the remote computer, the credentials are encrypted.

Backup Exec logon accounts can also be associated with backup data at the device level such as shares, databases, etc. If you need to edit the credentials, you can edit the Backup Exec logon account. Any changes are applied to the selected computers that use the Backup Exec logon account.

Backup Exec logon accounts are not user accounts. When you create a Backup Exec logon account, an entry for the account is entered into the Backup Exec database; no operating system accounts are created. If your user account credentials change, you must update the Backup Exec logon account with the new information. Backup Exec does not maintain a connection with the user account.

You can view, create, edit, replace, and delete Backup Exec logon accounts.

The following types of logon accounts are included in Backup Exec:

[Default Backup Exec logon account](#)

Backup Exec System Logon Account

Restricted logon accounts

Default Backup Exec logon account

The default Backup Exec logon account enables you to browse, make selections, or restore data. The first time you start Backup Exec, you must specify a default Backup Exec logon account using the Logon Account Wizard. You can select an existing Backup Exec logon account or create a new one.

You can create multiple Backup Exec logon accounts; however, each Backup Exec user can have only one default Backup Exec logon account.

Your default Backup Exec logon account enables you to perform the following:

- Browse data and make backup selections. Your default Backup Exec logon account enables you to browse local and remote computers when you create backup jobs provided your default account has sufficient rights.
You can select a different Backup Exec logon account when you make selections for backup. If your default logon account does not have rights, the **Logon Account Selection** dialog box appears and lets you create or select a different Backup Exec logon account.
See [“Changing your default Backup Exec logon account”](#) on page 768.
See [“Requirements for using the SQL Agent”](#) on page 1121.
See [“Requirements for accessing Exchange mailboxes”](#) on page 1175.
- Restore. You can assign Backup Exec logon accounts to computers when you create restore jobs. The default Backup Exec logon account is used unless you choose a different Backup Exec logon account when you create the restore job.

Backup Exec System Logon Account

The Backup Exec System Logon Account (SLA) is created when you install Backup Exec. When the SLA is created, the user name and password match the credentials that were provided during install for the Backup Exec Services credentials. The owner of the SLA is the user that installed Backup Exec. It is a common account, by default. Common accounts are the shared accounts that all users can access.

See [“Creating a new Backup Exec System Logon Account”](#) on page 768.

The Backup Exec System Logon Account may have access to most or all of your data since it contains the Backup Exec Services credentials. If you want to make Backup Exec more secure, you can change the SLA to be a restricted account. You can also delete it after making another logon account the default. However, if you delete the SLA, the jobs in which it is used may fail. If the SLA is deleted, you can re-create it using the **Logon Account Management** dialog box.

The SLA is used for the following tasks and jobs:

- Jobs that were migrated from a previous version of Backup Exec
- Duplicate backup data jobs
- Command Line Applet (bemcli.exe)

Restricted logon accounts

Backup Exec logon accounts can be common or restricted. When you create a Backup Exec logon account, you can designate it as a restricted account. To use a restricted logon account, you must be the owner of the logon account or you must know the password for the logon account. The person who created the logon account is the owner. If you authorize only a few people to back up or restore data, you can make the logon account a restricted logon account.

The main reasons to restrict a logon account are as follows:

- To help you limit access to the computers available for backup.
- To help you limit the computers to which you can restore.

When you use a restricted logon account to select the data for a job, the logon account information is saved with the selection list. Anyone who tries to edit the job must provide the password to the restricted logon account. Backup Exec loads the selections for that job only when the password for the restricted logon account is provided.

Some features available in the Backup Exec global settings let you set up a Logon Account. For example, Simplified Disaster Recovery where a Logon Account must be specified. To access global settings, click the **Backup Exec button > Configuration and Settings > Backup Exec Settings**. By default, the Backup Exec System Logon Account is assigned for a feature. Since these settings are global to Backup Exec, the assigned Logon Account is accessible to all logged in users of Backup Exec. As a security best practice, review the Logon Accounts in the global settings that are assigned to the features. It is recommended that you create and assign Logon Accounts with only minimal privileges required for that feature and select **Common** as the type of account that can be shared across all users of Backup Exec.

See [“Creating a Backup Exec logon account”](#) on page 762.

See [“Editing a Backup Exec logon account”](#) on page 763.

See [“Changing the password for a Backup Exec logon account”](#) on page 765.

See [“Replacing a Backup Exec logon account”](#) on page 766.

See [“Deleting a Backup Exec logon account”](#) on page 766.

See [“Copying logon account information to another Backup Exec server”](#) on page 769.

Creating a Backup Exec logon account

You can create Backup Exec logon accounts using the Logon Account Wizard, which guides you through the creation of a Backup Exec logon account, or by using the **Logon Account Management** dialog box. You can enter Backup Exec logon account property information when you create the Backup Exec logon account. However, Backup Exec assigns the Backup Exec logon account owner to the user name you used to log on to Backup Exec. The owner of the Backup Exec logon account cannot be modified.

This topic includes the following information:

[To create a Backup Exec logon account using the Logon Account Wizard](#)

[To create a Backup Exec logon account manually](#)

To create a Backup Exec logon account using the Logon Account Wizard

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Logon Account Wizard**.
- 3 Click **Add a new logon account**, and then click **Next**.
- 4 Type a user name and password.
- 5 Click **Next**.
- 6 In the **Logon account name** field, type the unique name for the Backup Exec logon account.
- 7 Under **Make this account**, select whether you want the account to be a common logon account or a restricted logon account.

Common logon accounts are the shared accounts that all users can access. Restricted logon accounts can only be used by the owner of the logon account or by those who know the password.

- 8 If you want to make this the default logon account that is used to browse, make selections, and restore data on your local computers and remote computers, select **The default logon account**.
- 9 Click **Next**.
- 10 Review the options that you selected, and then click **Finish** to create the logon account.

To create a Backup Exec logon account manually

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Manage Logon Accounts**.
- 3 Click **Add**.
- 4 In the **User name** field, type the fully qualified user name for the Backup Exec logon account.

For example, type "DOMAIN\Administrator".

The user name is provided when you attempt to connect to a computer. The user name is not case-sensitive for the computers that are accessed.

- 5 In the **Password** field, type the password for the account.
The password you enter is encrypted for security. You can leave this field blank if this Backup Exec logon account does not need a password.
- 6 In the **Confirm password** field, type the password again to verify it.
- 7 In the **Account name** field, type the unique name for the Backup Exec logon account.

- 8 In the **Notes** field, type any optional notes to explain how the Backup Exec logon account is used.
- 9 Select **This is a restricted logon account** if you want the Backup Exec logon account to be used only by the owner of the logon account and those who know the password.

If this option is not selected, the Backup Exec logon account is created as a common account. Common accounts are the shared accounts that all users can access.

- 10 Select **This is my default account** to make this account your default Backup Exec logon account, which is used to browse, make selections, or restore data on your local computers and remote computers.

See ["Backup Exec logon accounts"](#) on page 759.

Editing a Backup Exec logon account

When you edit a Backup Exec logon account, the changes are automatically applied to all the content that uses the Backup Exec logon account. Changes made to a Backup Exec logon account are applied immediately. You do not have to restart your system for the changes to take effect.

You can edit the following properties for a Backup Exec logon account:

- Type (restricted, common, or default)

- Account name
- Password
- User name
- Notes

To edit a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to change, and then click **Edit**.

If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

4 Edit any of the following properties:

User name	Type the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a computer. The user name you enter is not case sensitive for the computers that are accessed.
Change Password	Click this option to change the password for the account. The password you enter is encrypted for security.
Account name	Type a unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Type any optional notes indicating how the Backup Exec logon account is used.
This is a restricted logon account	Select this option to make this Backup Exec logon account a restricted logon account. Restricted logon accounts can be used only by the owner of the logon account and those who know the password. If this option is not selected, the Backup Exec logon account is a common account. Common accounts are the shared accounts that all users can access.
This is my default account	Select this option to make this account your default Backup Exec logon account. Your default account is used to browse, make selections, or restore data on your local computers and remote computers.

5 On the **Edit Logon Credentials** dialog box, click **OK**.

6 On the **Logon Account Management** dialog box, click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Changing the password for a Backup Exec logon account

You can change a Backup Exec logon account password using the following steps. Changes made to a Backup Exec logon account password are applied immediately.

To change the password for a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.

- 3 Select the Backup Exec logon account that you want to change, and then click **Edit**.

If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

- 4 Click **Change Password**.
- 5 In the **Password** field, type a new password.
- 6 In the **Confirm** field, re-type the password, and then click **OK**.
- 7 On the **Edit Logon Credentials** dialog box, click **OK**.
- 8 On the **Logon Account Management** dialog box, click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Replacing a Backup Exec logon account

You can replace a Backup Exec logon account within all existing jobs. The data in existing jobs that use the Backup Exec logon account will be updated to use the new Backup Exec logon account. If the new Backup Exec logon account is restricted, you must provide the password.

To replace a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to replace, and then click **Replace**.
- 4 On the **Replace Logon Account** dialog box, select the Backup Exec logon account with which you want to replace the selected Backup Exec logon account.

If the Backup Exec logon account is restricted and you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can select the account.

- 5 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Deleting a Backup Exec logon account

If you no longer need a Backup Exec logon account, you can delete it.

You cannot delete a Backup Exec logon account in the following situations:

- It is referenced by a job.
- It is owned by a user who is logged on to the Backup Exec server.
- It is set as the default Backup Exec logon account of a user who is logged on to the Backup Exec server.

If a logon account is used in any of these situations, you must replace it with a different logon account before you can delete it.

To delete a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to delete, and then click **Delete**.
- 4 Do either of the following:

If the logon account is not referenced by any backup jobs

Click **Yes** to confirm the deletion.

If the logon account is referenced by backup jobs

Complete the following steps to replace the logon account with a different logon account in any referenced backup jobs.

- On the **Delete Logon Account** dialog box, click **Replace**.
- On the **Replace Logon Account** dialog box, select the Backup Exec logon account with which you want to replace the selected Backup Exec logon account.

If the Backup Exec logon account is restricted and you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can select the account.

The logon account that you select here replaces the previous logon account in all existing jobs and selection lists.

- 5 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

See [“Replacing a Backup Exec logon account”](#) on page 766.

Changing your default Backup Exec logon account

You can change your default Backup Exec logon account that enables you to browse, make selections, or restore data.

To change your default Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account that you want to use as your default Backup Exec logon account, and then do one of the following:
 - Click **Set as Default**.
 - Click **Edit**, select **This is my default account**, and then click **OK**.
- 4 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Creating a new Backup Exec System Logon Account

The Backup Exec System Logon Account enables you to perform several operations. If you delete the Backup Exec System Logon Account, you should create a new one that enables you to perform the specified operations.

To create a new Backup Exec System Logon Account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Click **System Account**.

4 Complete the following options:

User name	Type the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a computer. The user name you enter is not case-sensitive for the computers that are accessed.
Change Password	Click this option to change the password for the account. The password you enter is encrypted for security.
Account name	Type a unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Type any optional notes indicating how the Backup Exec logon account is used.
This is a restricted logon account	Select this option to make this Backup Exec logon account a restricted logon account. Restricted logon accounts can be used only by the owner of the logon account and those who know the password. If this option is not selected, the Backup Exec logon account is a common account. Common accounts are the shared accounts that all users can access.
This is my default account	Select this option to make this account your default Backup Exec logon account. Your default account is used to browse, make selections, or restore data on your local computers and remote computers.

5 Click **OK** to create the system logon account.

See [“Backup Exec logon accounts”](#) on page 759.

Copying logon account information to another Backup Exec server

You can copy logon account information from one Backup Exec server to a different Backup Exec server.

To copy logon account information to another Backup Exec server

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.

- 3 Select the logon account that you want to copy, and then click **Copy to Servers**.
 - 4 If prompted, enter the password for the logon account that you selected.
 - 5 Do one of the following:
 - To add individual servers manually, in the **Server Name** field, enter the name of the Backup Exec server that you want to copy the logon account information to, and then click **Add**
 - To add several servers from a list, click **Import List**, and then browse to the list of server names.
 - 6 If you want to overwrite a logon account with the same name on the destination Backup Exec server, check **Overwrite logon account if one with this description already exists on the destination server**.
 - 7 Click **OK**.
- See [“Backup Exec logon accounts”](#) on page 759.

Testing logon accounts

You can test Backup Exec logon accounts to ensure that they can access your backup sources before you run jobs. Testing your logon accounts before you attempt to run jobs can help prevent failures and save you time.

To test logon accounts

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Click **Test**.
- 4 In the **Logon Account** field, select the logon account that you want to test.
- 5 In the **Server** field, select the server that you want to access with the logon account.
- 6 Click **Test**.

Backup Exec checks to make sure whether the logon account can access the server that you selected.
- 7 When you are finished testing logon accounts, click **Close**.
- 8 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 759.

Starting and stopping Backup Exec services

You can use the Backup Exec Services Manager to start, stop, and restart Backup Exec services.

To start or stop Backup Exec services

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 Select the server for which you want to start or stop services.
- 3 Do any of the following:

To start all services for the selected server Click **Start all services**.

To stop all services for the selected server Click **Stop all services**.

To restart all services for the selected server Click **Restart all services**.

- 4 Click **OK**.

See [“Changing the credentials for a service account”](#) on page 771.

See [“Changing startup options for Backup Exec services”](#) on page 773.

Changing the credentials for a service account

On the Backup Exec server, all Backup Exec services run in the context of a user account that is configured for the Backup Exec system services.

Note: The Backup Exec service account and the Backup Exec system logon account are set to the same user name when Backup Exec is installed. If you need to change the user name for the service account or if the service account is no longer used, then you should also change the Backup Exec system logon account to use new credentials.

If this computer is in a domain, enter a Domain Administrators account, or an equivalent account that is part of the Domain Admins group. In the Domain list, select or enter the Domain name.

If this computer is in a workgroup, enter an Administrators account, or an equivalent account that is part of the Administrators group on the computer. In the Domain list, select or enter the computer name.

The account that you designate for Backup Exec services, whether it is a new account or an existing user account, is assigned the following rights:

- Authenticate as any user and gain access to resources under any user identity.
- Create a token object, which can then be used to access any local resources.
- Log on as a service.
- Administrative rights (provides complete and unrestricted rights to the computer).
- Backup operator rights (provides rights to restore files and directories).
- Manage auditing and security log.

See [“Required user rights for backup jobs”](#) on page 150.

Due to security implementations in Microsoft Small Business Server, the service account must be Administrator.

To change the credentials for a service account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 On the **Backup Exec Services Manager** dialog box, select the appropriate server, and then select the service for which you want to change the service account.
- 3 Click **Edit credentials**.
- 4 Enter the credentials for the existing service logon account and click **Verify**.
After the authentication, you can change the service account credentials.
- 5 Select the **Change service account credentials** check box.
- 6 Complete the following fields:

New user name	Type the new user name for the service account that you want to change.
New password	Type the new password for the service account that you want to change.
Confirm password	Type the new password again to confirm it.
Grant required rights to the service account	Select this option to grant the service account the proper system service rights.

7 Click **OK**.

8 Click **Close**.

See [“Starting and stopping Backup Exec services”](#) on page 771.

See [“Changing startup options for Backup Exec services”](#) on page 773.

Changing startup options for Backup Exec services

You can change startup options for Backup Exec services. Each individual service can be configured to start automatically or manually. Or you can disable a service entirely.

Services that are configured for automatic startup automatically start when the server starts. Services that are configured for manual startup do not start automatically. You must manually start services that are configured for manual startup. You can start, stop, or restart services in the Backup Exec Services Manager.

To change service startup options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 On the **Backup Exec Services Manager** dialog box, select the appropriate server, and then select the service for which you want to change startup options.
- 3 Click **Edit credentials**.
- 4 Check the **Change startup options** check box.
- 5 Select from the following options:

Automatic

Select this option to automatically start the service account at system startup.

Manual

Select this option to prevent the service account from automatically starting at system startup. If you select this option, you must start the service account manually.

Disabled

Select this option to disable the service account at system startup.

6 Click **OK**.

7 Click **Close**.

See [“Starting and stopping Backup Exec services”](#) on page 771.

See [“Changing the credentials for a service account”](#) on page 771.

Configuring audit logs

You can use audit logs to examine and review information about the operations that have been performed in Backup Exec. The audit log displays the date and time of the activity, who performed it, what the activity was, and a description of the activity.

Audit logs can be configured to display information about the activities that occur for all or any of the following:

- Alerts
- Audit logs
- Backup set retention
- Devices and media
- Encryption keys
- Error-handling rules
- Install
- Jobs
- Logon accounts
- Server configuration

You can delete the audit logs as part of the Backup Exec database maintenance, and you can save the audit log to a text file. Any changes that are made to the audit log, such as when database maintenance occurs, can also be displayed in the audit log.

To configure audit logs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 Click **Configure Logging**.

- 3 On the **Audit Log Configuration** dialog box, select the check box of the category that you want to display in the audit log.

Expand the category by clicking the arrow to the left of the category. Select the operations that you want to display for the category.

Clear the check box of any item or operation that you do not want to display.
 - 4 Click **OK**.
- See [“Viewing the audit log”](#) on page 775.
- See [“Removing entries from the audit log”](#) on page 775.
- See [“Saving an audit log to a text file”](#) on page 776.

Viewing the audit log

You can view audit logs to see when changes were made in Backup Exec and which users made the changes.

To view the audit log

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 In the **Select category to view** field, select the category for which you want to view audit information.
- 3 Use the scroll bar at the bottom of the Audit Logs window to view the whole entry, or double-click the entry to display the same information in an easy-to-read Audit Log Record.

See [“Configuring audit logs”](#) on page 774.

Removing entries from the audit log

You can remove the entries for all categories or for a selected category.

To remove entries from the audit log

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 In the **Select category to view** field, select the category for which you want to view audit information.
- 3 Click **Clear Category Log** to remove all entries from an audit log category.

If you select specific categories, only the logs that are generated for the selected categories are cleared when you click **Clear Category Log**.

See [“Configuring audit logs”](#) on page 774.

Saving an audit log to a text file

You can save the audit log as a text file.

To save the audit log to a text file

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 Click **Save Log to File** to specify a file name and location to save the audit log entries.

See [“Configuring audit logs”](#) on page 774.

Copying configuration settings to another Backup Exec server

If you have the Central Admin Server feature (CAS), you can copy configuration settings and logon information from one Backup Exec server to another. This copy ability lets you quickly set up a group of Backup Exec servers with the same configuration or logon settings.

Note: To copy configuration settings and logon information to other Backup Exec servers, you must install the **Copy Server Configurations** feature.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

To copy configuration settings to another Backup Exec server

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Copy Settings to Other Servers**.

- 2 Select any of the following options:

Default schedule	Select this option to copy the default schedule settings from this Backup Exec server to another Backup Exec server.
Error-handling rules	Select this option to copy error-handling rules from this Backup Exec server to another Backup Exec server.
Alert configuration	Select this option to copy the alert configuration from this Backup Exec server to another Backup Exec server.

- 3 Select the Backup Exec server or servers to which you want to copy the selected settings.

If the destination server is not in the list, do one of the following:

- To add a server manually, click **Add**, and then select the server or servers that you want to add to the list.
- To import a list of servers, click **Import List**, browse to select the list, and then click **Open**.

- 4 Click **OK**.

See [“About the Central Admin Server feature”](#) on page 1321.

Viewing server properties

You can view properties for the local Backup Exec server or any other server that you monitor with Backup Exec.

You can view the following properties for the local Backup Exec server:

- Server name
- Server description
- Server status
- Version and license information
- Date and time zone information

- Operating system information
- Memory and page file information

Additionally, you can view the following properties for any server that you monitor with Backup Exec:

- Server name
- Server description
- Operating system information
- Backup Exec version and license information

You can also view the following properties for a Microsoft 365 Tenant

- Type
- Tenant ID
- Certificate Expiry
- Azure Active Directory application details

If you have the Central Admin Server feature (CAS), you can also view information about the Backup Exec database, device and media database, and catalog database.

See [“Viewing the settings for a central administration server”](#) on page 1375.

To view server properties

- ◆ Complete either of the following as necessary:

To view the local Backup Exec server's properties

- Click the Backup Exec button, select **Configuration and Settings**, and then click **Local server properties**.
- When you are finished viewing the local server properties, click **OK**.

To view any other server's properties

- On the **Backup and Restore** tab, double-click the server whose properties you want to view.
- In the left pane, click **Properties**.

Configuring default backup settings

You can choose the processing method to use for Hyper-V backups. You set this option for the Backup Exec server, so the method that you choose is applied to all backups of Hyper-V virtual machines that the Backup Exec server protects.

Three processing options are available:

- The Resilient Change Tracking (RCT) method. This method tracks the changes for virtual machines that should be backed up. The RCT method provides better resiliency than the Standard and Faster processing methods. Backups are faster for larger virtual machines compared to other 2 methods. RCT does not make use of VSS infrastructure on the Hyper-V host. This method is available only for Hyper-V servers that run Windows 2016 and later.
- The standard processing method. This method reads the whole virtual disk and identifies the changes that should be backed up. The changed blocks that are identified are then backed up.
- The faster processing method. This method is faster than the standard processing method because it writes all changes to a new differencing disk and then backs up only the differencing disk. This method saves time because the entire disk does not have to be read.

Table 17-14 Difference between the Resilient Change Tracking, standard processing, and the faster processing methods

Processing method	Supported Hyper-V servers	Backup type	Disk storage space	Impacts on system performance
Resilient Change Tracking (RCT) method	Windows Server 2016 and later.	■ Differential ■ incremental	No extra space is required.	Does not impact performance
Standard processing method	Windows Server 2012 and later.	■ Differential ■ Differential with incremental	No extra space is required.	Does not impact performance

Table 17-14 Difference between the Resilient Change Tracking, standard processing, and the faster processing methods (*continued*)

Processing method	Supported Hyper-V servers	Backup type	Disk storage space	Impacts on system performance
Faster processing method	Windows Server 2012 and later.	■ Incremental ■ Full ■ Differential backups run as incremental backups and the job completes with a status of "Success with exceptions".	Requires extra space on the Hyper-V host even after the backup job is complete. The extra space required depends on how long the checkpoint is not merged back into the parent disk and the number of writes that have happened before the merge. The space consumed for each virtual machine might not be significant but if there are many virtual machines in your environment, the checkpoints consumes a significant amount of space. However, you can use this method if you take frequent backups and disk space is not a constraint.	May slow down system performance because a checkpoint is always present for each virtual machine backed up using this method.

How upgrades affect Hyper-V backup settings

The following notes provide information for upgrades:

- When you upgrade from Backup Exec 15 Feature Pack 3 and later, the existing Hyper-V backup setting does not change. Existing and new backup jobs use the setting that was configured in Feature Pack 3, unless you change it.
- When you upgrade from Backup Exec 15 Feature Pack 2 or earlier, the default Hyper-V backup setting for existing jobs is **Use the standard processing method**. Existing and new backup jobs use this setting unless you change it. In this scenario, the Resilient Change Tracking (RCT) method is disabled.
- On the Backup Exec server, it is always recommended that you select the Resilient Change Tracking (RCT) method. If RCT is selected then it is used wherever applicable. If RCT is not supported by a virtual machine, then one of the two methods, standard processing or faster processing is used.

To configure Hyper-V incremental or differential backup settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Virtual Machines**.

- 3 Select the processing method that you want to use for all Hyper-V incremental or differential backup jobs.

Use Resilient Change Tracking wherever applicable

Select this option if you want to run both incremental backups and differential backups. If RCT is selected then it is used wherever applicable and is the recommended method. If a virtual machine does not support RCT, the standard or faster processing method is used, based on your selection.

Note: This option is available only for Hyper-V servers that run Windows Server 2016 and later.

If you do not select the RCT method and if a Windows Server 2016 or later host is found, Backup Exec sends periodic alerts for you to enable the RCT method. If you do not want Backup Exec to display the alert messages, you can disable the alerts by editing the following registry key:

Registry Location:

```
HKEY_LOCAL_MACHINE\Software\Veritas\Backup
Exec For Windows\Backup Exec\Server
ValueName = SupressRCTAlert
```

When you set the value as **1**, the alerts are disabled.

Use the faster processing method

Select this method if you want incremental backups jobs to be processed as quickly as possible, and do not want to perform any differential backups. This option does not support differential backups. If you select this option, all Hyper-V differential backups are processed as incremental backups.

Note: This option is available only for Hyper-V servers that run Windows 2012 and later. For all supported previous versions of Windows, the standard processing method must be used.

Use the standard processing method

Select this option if you want to run both incremental backups and differential backups.

Note: When you change to a different backup processing method, the next job runs as a full backup instead of an incremental or a differential backup.

Note: In a CAS environment, if the processing method is not the same on the central administration server and the managed Backup Exec server, the method that is set for the managed Backup Exec server is used when a job is delegated from the central administration server.

- 4 Click **OK**.

Changing virtual machine validation settings for VMware and Hyper-V

In the **Virtual Machine Validation Settings**, you can change the timeout settings of a virtual machine for VMware and Hyper-V. After you create the validation job for a virtual machine, before the job runs, you can change the maximum time that it takes for a virtual machine to boot. This is a global setting that is applicable to all validation jobs. If the virtual machine does not boot within the selected time, the validation job fails.

To change virtual machine validation settings for VMware and Hyper-V

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Virtual Machines**.
- 3 In the **Select the maximum boot time for a virtual machine** option, select the time in minutes.

By default, 5 minutes is selected. You can select from 1 to 60 minutes.

Changing malware scan options

You can configure malware scan options in Backup Exec. The malware scan options are global options that can affect all Backup Exec malware scan jobs.

To change malware scan options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Malware Scan**.
- 3 Complete the following options:

Malware Scan Engine Name	Select scan engine that you want to use to perform malware scan. The available options are: ■ MS Defender ■ Symantec Protection Engine (SPE) By default, the MS Defender will be used as the scan engine for malware scan. You may change it to Symantec Protection Engine (SPE).
Allow the malware scan job to use the maximum percentage of available CPU	Select the setting to use the maximum percentage of available CPU while running a malware scan job to work with other processes in Backup Exec. Note: The option to use maximum percentage of CPU is not available for Symantec Protection Engine(SPE), so this option is disabled.
Number of parallel threads	Select the number of streams you want to use for the malware scan jobs. The maximum number of streams cannot exceed 10.

4 Click **OK**.

Reports

This chapter includes the following topics:

- [Reports in Backup Exec](#)
- [Running a report now](#)
- [Scheduling a report](#)
- [Creating a custom report](#)
- [Saving a report](#)
- [Printing a report from the Backup Exec Report Viewer](#)
- [Viewing completed reports](#)
- [Editing a report](#)
- [Re-running a completed report](#)
- [Deleting a report](#)
- [Setting defaults for standard and custom reports](#)
- [Viewing report properties](#)
- [List of Backup Exec standard reports](#)

Reports in Backup Exec

Backup Exec includes more than 40 standard reports that can provide detailed information about the alerts, devices, media, and jobs in your Backup Exec environment. In addition, Backup Exec provides the ability to create custom reports so you can create reports to fit your specific environment and needs. Both standard reports and custom reports can be run on demand at any time,

can be scheduled to run at a specific time, or can be scheduled to run on a recurring schedule. When you schedule a report job, you can set up email notification so that you or anyone else can be notified when the report job completes, and you can include a copy of the report in the email

Reports can be viewed and printed in the following formats:

- PDF
- HTML
- XML
- Microsoft Excel (XLS)
- Comma-separated Value (CSV)

Reports are grouped on the **Reports** tab by category. For example, reports relating to devices are grouped in the **Devices** report group. To see the names and descriptions of the standard reports for a group, click the name of the appropriate report group.

Note: The **See Completed** report group only includes the scheduled reports that have completed. Reports that were run immediately are deleted when you close the report viewer.

For information about the best practices to use Backup Exec reports, refer to the *Backup Exec Best Practices*.

Requirements for reports

Before you attempt to run either standard reports or custom reports, review the following requirements:

- To properly format integrated Backup Exec reports, you must configure a default printer using the Windows Control Panel Printers applet. This is required even if you do not have a printer attached to your system. For information on configuring a printer by using the Windows Control Panel Printers applet, see your Microsoft Windows documentation.
- To run reports across multiple Backup Exec servers, you must install the Backup Exec Enterprise Server feature, even if you do not have a shared storage environment.
- To view reports in PDF format, ensure that the latest version of Adobe Acrobat Reader is installed.

Reports and the Job Monitor

Reports cannot be monitored from the **Job Monitor**. Reports that are running, are scheduled to run, or are completed do not appear on the **Job Monitor**. All report operations are displayed on the **Reports** tab.

See [“Running a report now”](#) on page 787.

See [“List of Backup Exec standard reports”](#) on page 803.

Running a report now

When you run a report, you can specify the criteria that is used to determine the items that will be included in the report. The settings that are available for you to select depend on the type of data that can be included in the report. After the report is generated, only the items that match the criteria appear in the report.

If you do not want to run the report immediately, you should use the **New Scheduled Report** option instead.

See [“Scheduling a report”](#) on page 787.

To run a report

- 1 On the **Reports** tab, under **Report Groups**, click the report group that contains the report that you want to run now.
- 2 Right-click the report you want to run, and then click **Run Now**.
- 3 If the **Run Report Now - <report_name>** properties page appears, select the appropriate settings for the data that you want to include in the report, and then click **OK**.
- 4 After you have finished viewing the report, click **Close**.

Backup Exec automatically deletes the report when you close the Report Viewer.

Scheduling a report

You can schedule a report to run once at a specific time or multiple times on a recurring schedule. Scheduled reports are listed in the **See Upcoming** report group on the **Reports** tab. You can edit or delete scheduled reports.

To schedule a report

- 1 On the **Reports** tab, under **Report Groups**, click the report group that contains the report that you want to schedule.
- 2 Right-click the report name, and then click **Schedule Report**.

- 3 Type a name for the report.
- 4 If any of the following options appear in the left pane, click the option name, and then enter the criteria for the data that you want to include on the report. Note that some reports do not include any of the following options.
 - **Validation Status**
 - **Backed Up Servers**
 - **Ranges**
 - **Audit Log Category**
 - **Job Status**
 - **Media Sets**
 - **Vault**
 - **Anomaly Status**
 - **Anomaly Severity**
- 5 (Optional) If you want to send an email to yourself or someone else when the report is completed, do the following:
 - In the left pane, click **Notification**.
 - Select an existing recipient or click **Manage Recipients** to add a new recipient.
 - If you want to include a copy of the report in the email notification, check **Include the report in email notifications**.

- 6 In the left pane, click **Schedule**, and then select one of the following scheduling options:

Recurrence	Lets you schedule the job to run multiple times. You can set the recurrence pattern so that the job runs
Run now with no recurring schedule	Lets you run the job immediately.
Run on	Lets you set the date and time to run the report.
Create without a schedule	Lets you create and save the report, but not schedule or run the report at this time. If you select this option, you can then use an external scheduling tool to schedule the job to run at a specific time, or you can manually run the job when you are ready. Even though the job is not scheduled, the report is listed in the See Upcoming report group on the Report tab.

- 7 Click **OK**.

The scheduled report is saved in the **See Upcoming** report group. After the report runs, it is saved in the **See Completed** report group.

See [“Viewing completed reports”](#) on page 800.

Creating a custom report

You can create reports that contain information to meet the specific requirements of your organization. You choose the data to include in the report, and then determine how the data is filtered, sorted, and grouped. In addition, you can set up a pie graph or a bar graph to graphically represent the report data.

Filters let you customize reports to include only the information that meets specific criteria. For example, you can use filters to find the jobs that contain a specific word, the alerts that occurred on a specific day, or the media that are in a specific location. You use filter criteria to create filter expressions. You can use one or multiple filter expressions. A filter expression consists of a field name, an operator, and a value.

The following example filter expression finds all alerts for errors:

Table 18-1 Filter expression for finding alerts for errors

Filter type	Data
Field name	Alert Type
Operator	= (Equal)
Value	Errors

If you want the report to include only the alerts for errors that occurred on a specific day, add another filter expression for the date and time, as shown in the example below:

Table 18-2 Filter expression for finding alerts on a specific day

Filter type	Data
Field name	Date Entered
Operator	=(Equal)
Value	06/03/2014 <time>

Grouping fields creates sections on the report. For example, if you group by Backup Exec server, Backup Exec creates a section for each Backup Exec server that matches the filter criteria. Under each Backup Exec server's section, the report displays the data that corresponds to the remaining fields that you selected for the report.

You can sort a custom report by up to three of the fields that you have chosen for the report. When you sort on fields, Backup Exec arranges all of the data that matches the sort criteria together in the report. For example, if you sort on the Backup Exec server field in ascending order, all data for Backup Exec server A displays first, followed by all data for Backup Exec server B, and so on.

To create a custom report

- 1 On the **Reports** tab, click **New Custom Report**.
- 2 On the **Custom Report** dialog box, type a name and description for the report.

- 3 If you do not want this report to include the default header and footer settings, uncheck **Use header and footer settings specified in Backup Exec Settings**.

The default header and footer settings can include a customized logo, a custom color for the banner, and text for the footer. These items are set in the default Backup Exec settings.

See [“Setting defaults for standard and custom reports”](#) on page 801.

- 4 In the left pane, click **Field Selection**.
- 5 In the **Category** box, select a group for which you want to create a report.
- 6 For additional field selections, click **Show advanced fields**.
- 7 In the **Available fields** list, select the fields that you want to include on the report, and then click the Right arrow (>) button to move the fields to the **Fields selected for the report** list.
- 8 After you have moved all of the fields that you want to include on the report to the **Fields selected for the report** list, arrange the order in which you want the fields to appear.

The fields appear on the report in the order in which they are listed in the **Fields selected for the report** list, with the first field appearing as the first field on the left of the report. To move a field, select it from the **Fields selected for the report** list, and then click **Move Up** or **Move Down** to move it to the appropriate location in the list.

- 9 To adjust the width of the column for a field, do the following in the order listed:
 - Click the field name in the **Fields selected for the report** list.
 - In the **Column width** field, type the new width.
 - Click **Set**.
- 10 (Optional) If you want to filter the data on the report, do the following:
 - In the left pane, click **Filters**.
 - In the **Field name** list, select the field on which you want to filter data.
 - In the **Operator** list, select the operator that you want to use for this filter.
 - In the **Value** field, type or select the specific data that you want to include on the report.
 - Click **Add**.
 - To combine sets of filter expressions, do any of the following:

To combine two filter expressions so that both expressions must be true for the result to be true

Click **AND**.

For example, to find all backup jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup

After you set up the expressions, do the following:

- Click **AND** to combine the two expressions.

The combined expression is:

Status = Failed AND Type = Backup

To combine two filter expressions so that one of the expressions must be true for the result to be true

Click **OR**.

For example, to find jobs that either failed or were canceled, add the following expressions:

- Status = Failed
- Status = Canceled

After you set up the expressions, do the following:

- Click **OR** to combine Status = Failed with Status = Canceled.

The combined expression is:

Status = Failed OR Status = Canceled

To combine two filter expressions into a single expression

Click **() +**

For example, to find backup jobs and restore jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup
- Type = Restore

After you set up the expressions, do the following:

- Use **OR** to combine Type = Backup with Type = Restore.
- Press and hold **Ctrl** while you click Type = Backup and Type = Restore.
- Click **() +** to combine Type = Backup with Type = Restore.
- Use **AND** to combine Status = Failed with (Type = Backup OR Type = Restore).

The combined expression is:

Status = Failed AND (Type = Backup OR Type = Restore)

To separate two filter expressions that were combined into a single expression

Click () -

For example, if you used () + to combine Type = Backup with Type = Restore, it is displayed on the **Filters** dialog box as follows:

(Type = Backup OR Type = Restore)

To make the combined expression into two individual expressions, do the following:

- Press and hold Ctrl while you click both Type = Backup and Type = Restore.
- Click () -

After you separate the expressions, they are displayed without the parentheses.

11 (Optional) If you want to organize the report into sections, do the following:

- In the left pane, click **Grouping**.
- In the **Group by** list, select the field that you want to use as a group.
- Click **Ascending** or **Descending**.
Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
- If you want to further group the data, select the fields in the **Then group by** lists, and then click **Ascending** or **Descending** for those fields.
A report must have at least one field that is not grouped. For example, if you select three fields to include on the report, you can group only two of the fields. If you group all of the fields, no data appears on the report because all of the data is listed in the group section titles. In addition, you must have at least four fields on the report to use all three grouping fields.

12 (Optional) If you want to sort the data on the report, do the following:

- In the left pane, click **Sort**.
- In the **Sort on** list, select the field on which you want to sort the data in the report.
- Select **Ascending** or **Descending**.
Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
- If you want to further sort the data, select the fields on which you want to sort in the **Then sort on** lists, and then click **Ascending** or **Descending** for those fields.

- 13 (Optional) If you want to add a pie graph or bar graph to the report, do the following:
 - In the left pane, click **Graph Options**.
 - In the **Graph type** list, select either **Bar** or **Pie**.
 - Select the types of data that you want to include on the graph.
- 14 (Optional) If you want to see what the report will look like before you save it, in the left pane, click **Preview**.
- 15 Click **OK** to save the custom report.

Adding or removing fields on a custom report

To change the data that appears on a custom report, add new fields or remove existing fields.

To add or remove fields on a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report that contains the fields you want to change, and then click **Edit**.
- 3 In the left pane, click **Field Selection**.
- 4 Do any of the following:

To add new fields to the report

- Select a category.
- Under **Available fields**, select the fields you want to add, and then click the right arrow (>) button.

To remove fields from the report

Under **Fields selected for the report**, click the field you want to remove, and then click the left arrow (<) button.

Changing filters for a custom report

Use the following steps to change the existing filters on a custom report.

To change filters for a custom report

- 1 On the **Reports** tab, under **Report groups**, click **Custom Reports**.
- 2 In the reports list, right-click the report that you want to change and then click **Edit**.

3 In the left pane, click **Filters**.

4 Create a filter by defining one or more filter expressions.

To add a new filter expression	Select a field name and operator, and then enter a value. Click Add .
To edit an existing filter expression	Do the following in the order listed. ■ Under Filter criteria, select the filter expression that you want to edit, and then click Edit. ■ Under Filter expression, edit the expression's values. ■ Under Filter expression, click Update.
To delete a filter expression	Under Filter criteria , select the filter expression that you want to delete, and then click Remove .

5 To combine sets of filter expressions, do any of the following:

To combine two filter expressions so that both expressions must be true for the result to be true	Click AND. For example, to find all backup jobs that failed, add the following expressions: ■ Status = Failed ■ Type = Backup After you set up the expressions, do the following: ■ Click AND to combine the two expressions. The combined expression is: Status = Failed AND Type = Backup
To combine two filter expressions so that one of the expressions must be true for the result to be true	Click OR. For example, to find jobs that either failed or were canceled, add the following expressions: ■ Status = Failed ■ Status = Canceled After you set up the expressions, do the following: ■ Click OR to combine Status = Failed with Status = Canceled. The combined expression is: Status = Failed OR Status = Canceled

To combine two filter expressions into a single expression

Click () +

For example, to find backup jobs and restore jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup
- Type = Restore

After you set up the expressions, do the following:

- Use OR to combine Type = Backup with Type = Restore.
- Press and hold Ctrl while you click Type = Backup and Type = Restore.
- Click () + to combine Type = Backup with Type = Restore.
- Use AND to combine Status = Failed with (Type = Backup OR Type = Restore).

The combined expression is:

Status = Failed AND (Type = Backup OR Type = Restore)

To separate two filter expressions that were combined into a single expression

Click () -

For example, if you used () + to combine Type = Backup with Type = Restore, it is displayed on the **Filters** dialog box as follows:

(Type = Backup OR Type = Restore)

To make the combined expression into two individual expressions, do the following:

- Press and hold Ctrl while you click both Type = Backup and Type = Restore.
- Click () -

After you separate the expressions, they are displayed without the parentheses.

6 Click **OK**.

Changing the way data is grouped or sorted in a custom report

Use the following steps to change the way data is grouped or sorted in a custom report.

To change the way data is grouped or sorted in a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 In the list of custom reports, right-click the report that you want to change, and then select **Edit**.

3 Do any of the following:

To change the field that is used as a group	■ In the left pane, click Grouping. ■ On the Group by or Then group by list that contains the field you want to change, click the down arrow, and then select the new field to use as a group.
To remove a group	■ In the left pane, click Grouping. ■ On the Group by or Then group by list that contains the group you want to remove, click the down arrow, and then click <None>.
To change the order in which data is grouped	Click Ascending or Descending .
To change the field that is used to sort the data	■ In the left pane, click Sorting. ■ On the Sort on or Then sort on list that contains the sorting option you want to change, click the down arrow, and then select the new field to use to sort the data.
To disable sorting	■ In the left pane, click Sorting. ■ On the Sort on or Then sort on list that contains the sorting option you want to remove, click the down arrow, and then click <None>.
To change the order in which data is sorted	Click Ascending or Descending .

4 Click **OK**.

Changing graph options in custom reports

You can include a pie graph or a bar graph in custom reports.

At least two fields must be selected on the **Field Selection** dialog box to create a pie graph, and at least three fields must be selected to create a bar graph.

To change graph options in custom reports

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report you want to edit, and select **Edit**.

- 3 In the **Custom Report** window, click **Graph Options**.
- 4 Change the graph title or select new fields to populate the graph.
- 5 Click **OK**.

Previewing custom reports

Use the preview feature to verify that you created a custom report correctly.

Note: You cannot preview custom reports from the Backup Exec Remote Administration Console.

To preview custom reports

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report that you want to preview, and then click **Edit**.
- 3 In the left pane, click **Preview**.

Copying a custom report

You can make one or more copies of a custom report. Each copy of the custom report resides in the Custom report group, along with the original custom report.

To copy a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click a custom report that you want to copy, and then click **Copy**.
- 3 Type a name for the report, and then click **OK**.

The copy of the custom report appears in the **Custom** report group.

Saving a report

Reports can be saved to any location that you choose on your hard drive or network in any of the following formats:

- HTML file (.htm)
- Adobe PDF file (.pdf)
- XML file (.xml)
- Comma-separated value file (.csv)
- Microsoft Office Excel workbook (.xls)

You can save a report that is currently displayed on the screen or that is in the **See Completed** report group.

To save a report

- 1 On the report, in the Report Viewer, click **Save As**.

To save a report that is currently displayed on the screen

On the report, in the Report Viewer, click **Save As**.

To save a report that is in the See Completed report group

- On the **Reports** tab, under **Report Groups**, click **See Completed**.
- Click **Save As**.
- Double-click a completed report that you want to save to a new location.

- 2 Enter the file name and location where you want to save the report.
- 3 In the **Save as type** box, select a format in which to save the report.

When you save a report in HTML format, both the HTML file and a .GIF image file are saved.

- 4 Click **Save**.

Printing a report from the Backup Exec Report Viewer

You can print reports from a locally-attached printer or a network printer. To print a report, the printer must be configured to print in the landscape mode.

The following printer settings ensure that the report prints correctly:

- On the **Layout** tab of the **Print** dialog box, under **Orientation**, **Landscape** should be selected. Note that you may need to select **Preferences** to access the **Layout** tab.
- To print all of the pages in a multiple page report, on the **Options** tab of the **Print** dialog box, the option **Print all linked documents** should be selected.
- To print all of the pages in a multiple-page report in the correct order, the first page of the report must be displayed on the screen before the print job is initiated.
- If the report does not print correctly, you may need to change the page setup options in Internet Explorer to remove the header and footer and reduce the margins.

To print a report from the Backup Exec Report Viewer

- 1 Run a report.
See [“Running a report now”](#) on page 787.
- 2 On the **Report Viewer**, click **Print**.
- 3 Select a printer from the Windows **Print** dialog box.
- 4 Click **Print**.

Viewing completed reports

After a scheduled report runs, it is saved in the **See Completed** report group.

To view completed reports

- 1 On the **Reports** tab, under **Report Groups**, click **See Completed**.
- 2 Double-click the report that you want to view.

Editing a report

Use the following steps to edit the properties of a scheduled standard report or a custom report before it runs. If the report that you want to edit has been run in a previous report job, the changes you make now may affect the appearance of the reports in job history. It is recommended that you copy the report and then edit the copy.

To edit a scheduled standard report or a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **See Upcoming**.

To edit a scheduled standard report

On the **Reports** tab, under **Report Groups**, click **See Upcoming**.

To edit a custom report

On the **Reports** tab, under **Report Groups**, click **Custom**.

- 2 Right-click a report you want to edit, and then click **Edit**.
- 3 Edit the report properties and then click **OK**.

Re-running a completed report

You can run the reports that appear in the **See Completed** report group multiple times.

To re-run a completed report

- 1 On the **Reports** tab, under **Report Groups**, click **See Completed**.
- 2 Right-click a report, and then click **Retry Report Now**.
Backup Exec creates and runs another iteration of the report.
- 3 To view the report again, double-click the new report.

Deleting a report

Reports that you create using the **Run now** option are automatically deleted after you view the report. However, custom reports, completed reports, and scheduled reports can be deleted at your convenience.

Note: Standard Backup Exec reports cannot be deleted.

To delete a report

- 1 On the **Reports** tab, under **Report Groups**, do one of the following:
 - To delete a custom report, click **Custom**.
 - To delete a scheduled report, click **See Upcoming**.
 - To delete a completed report, click **See Completed**.
- 2 Right-click the report that you want to delete, and then click **Delete**.
- 3 Click **Yes** to confirm that you want to delete the report.

Setting defaults for standard and custom reports

You can set Backup Exec to display all reports in either HTML or Adobe Portable Document Format (PDF). The default setting is HTML. The format that you select does not affect the format of the reports that are sent to users with the notification feature.

For custom reports, you can do the following:

- Include a logo in the header.
- Choose a color for the banner in the header.
When you choose a color for the banner, you can type the numbers that correspond to the colors (RGB values), or you can select the color from a chart.
- Include text in the footer.
- Include the time in the footer.

To set defaults for standard and custom reports

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Reports**.
- 3 Complete the appropriate options.

To select a default report format for standard and custom reports

Under **Report format**, select either **HTML** or **PDF**.

To limit the number of rows that can be included in standard and custom reports

Under **Report content**, in **Maximum number of rows to include in a report**, type the appropriate number.

To enable standard and custom reports to show all data that is generated, even if some of the data is duplicated

Under **Report content**, click **Show all rows**.

To enable standard and custom reports to show only unique data

Under **Report content**, click **Show distinct rows**.

To add a logo to the header section of custom reports

Under **Header**, check **Use custom image file**, and then enter the path to the image that you want to use in the Image file path field.

To customize the colors in the header section of custom reports

Under **Banner color**, either enter the numbers that correspond to the colors you want to use, or click **Colors** to select a color from a chart.

To include default text or the time in the footer of custom reports

Under **Footer**, in the **Text** field, type the default text that you want to appear on every custom report. If you want the time of the report to be included in the footer, verify that **Include time** is checked.

- 4 Click **OK**.

Viewing report properties

Report properties provide detailed information about each report, such as the file name, file size, and report creation date. The properties can be viewed, but not edited.

To view report properties

- 1 On the **Reports** tab, under **Report Groups**, select a report group.
- 2 Right-click a report for which you want to view properties, and then click **Properties**.
- 3 Click **OK** after you have finished viewing the properties.

List of Backup Exec standard reports

This section provides detailed information about each standard report that is available in Backup Exec. The data that appears on each report varies depending on the criteria you selected to include in the report.

The following standard reports are included in Backup Exec:

Table 18-3 Backup Exec standard reports

Report Name	Description
Alert History	Lists all alerts in the alert history chronologically, displaying the most recent alerts first See “Alert History report” on page 808.
Alert History by Backup Exec server	Lists all alerts in the alert history, grouped and filtered by Backup Exec server, displaying the most recent alerts first. See “Alert History By Backup Exec Server report” on page 808.
Anomaly Detection Summary	Lists all the anomalies detected by monitoring backup job parameters. See “Anomaly Detection Summary report” on page 809.
Audit Log	Lists the contents of the audit logs for selected servers for the specified time period. See “Audit Log report” on page 810.
Backup Job Success Rate	Lists the success rate for backup jobs run to back up selected servers. See “Backup Job Success Rate report” on page 811.

Table 18-3 Backup Exec standard reports (*continued*)

Report Name	Description
Backup Recommendations	Lists any recommendations that can help you to better manage your backups. See “Backup Recommendations report” on page 811.
Backup Resource Success Rate	Lists the success rate for backup jobs for specified past number of days for resources on selected servers. See “Backup Resource Success Rate report” on page 812.
Backup Sets by Media Set	Lists all backup sets by media set. See “Backup Sets by Media Set report” on page 812.
Backup Size by Resource	Lists the backup size for each resource job for up to seven previous runs and then computes the trailing average for up to seven previous runs for each job run. See “Backup Size By Resource report” on page 813.
Cloud Storage Summary	Displays the summary of the size of cloud storage that the backed up data uses on the Backup Exec server. See “Cloud Storage Summary report” on page 813.
Daily Device Utilization	Lists the percentage of the storage devices’ capacity that the Backup Exec server uses. See “Daily Device Utilization report” on page 814.
Deduplication Disk and Cloud Deduplication Device Summary	Displays a summary of the deduplication operations for local deduplication disk storage, shared deduplication disk storage, and cloud deduplication disk storage. See “Deduplication Disk and Cloud Deduplication Device Summary report” on page 815.
Deduplication summary	Displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec server. See “Deduplication Summary report” on page 816.

Table 18-3 Backup Exec standard reports (*continued*)

Report Name	Description
Device Summary	Lists the device usage and error summary for each selected Backup Exec server. See “Device Summary report” on page 817.
Disk Storage Summary	Displays disk-based usage statistics for Backup Exec server disk storage. See “Disk Storage Summary report” on page 817.
Error-Handling Rules	Lists all the defined error-handling rules. See “Error-Handling Rules report” on page 818.
Event Recipients	Lists all events that were received by each notification recipient. See “Event Recipients report” on page 819.
Failed Backup Jobs	Lists all the failed backup jobs, sorted by the resource server and time frame. See “Failed Backup Jobs report” on page 820.
Jobs Summary	Lists all the jobs that ran within the last 72 hours in chronological order. See “Jobs Summary report” on page 821.
Malware Scan Summary	Lists all the malware that were detected during the scan of the backup sets for the protected resources. See “Malware Scan Summary report” on page 821.
Managed Backup Exec Servers	Lists the status and configuration for all Backup Exec servers that are managed by a central administration server. See “Managed Backup Exec Servers report” on page 822.
Media Audit	Lists the recent media configuration changes. See “Media Audit report” on page 824.
Media Errors	Lists the number of errors that occur on all media. See “Media Errors report” on page 825.

Table 18-3 Backup Exec standard reports (*continued*)

Report Name	Description
Media Required for Recovery	Lists the media that contain the backup sets for each system that is backed up on selected servers for the specified time period. This report can be inaccurate if media overwrite settings allow the media to be overwritten. See “Media Required for Recovery report” on page 825.
Media Summary	Lists all the media sets and media that are used by Backup Exec servers. The current location is given for each media. Also lists usage statistics for media and the location of media within Backup Exec media sets. See “Media Summary report” on page 826.
Media Vault Contents	Lists the media that are located in each media vault. See “Media Vault Contents report” on page 827.
Move Media to Vault	Lists all media that can be moved to a media vault. The listed media are not currently in a media vault and the media’s append period has expired. See “Move Media to Vault report” on page 827.
Operations Overview	Lists past and future operations data for user-set period. See “Operations Overview report” on page 828.
Overnight Summary	Lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were scheduled to run but did not run. Jobs are given a grace period of 24 hours before they are marked as past due. See “Overnight Summary report” on page 830.
Problem Files	Lists all the problem files that are reported for jobs. The files are grouped by day and resource. See “Problem Files report” on page 831.
Recently Written Media	Lists all media that have been modified in the last 24 hours. See “Recently Written Media report” on page 832.

Table 18-3 Backup Exec standard reports (*continued*)

Report Name	Description
Recovery Ready Validation Summary	Lists the backup sets for which the Validate Virtual Machine for Recovery jobs are run. See “Recovery Ready Validation Summary” on page 832.
Resource Protected Recently	Lists all job detail statistics and exceptions that occurred on a Backup Exec server for which you run this report. See “Resource Protected Recently report” on page 833.
Resource Risk Assessment	Lists job information for resources on which the last backup job run on the resource failed. The data is filtered by resource server. See “Resource Risk Assessment report” on page 834.
Restore Set Details by Resource	Lists all restore sets that ran within the last 72 hours. The sets are grouped by the server and resource. See “Restore Set Details by Resource report” on page 835.
Retrieve Media from Vault	Lists all reusable media currently in the specified vault. See “Retrieve Media from Vault report” on page 836.
Robotic Library Inventory	Lists the contents of slots in robotic libraries that are attached to Backup Exec servers. Usage statistics are provided for each piece of media. See “Robotic Library Inventory report” on page 836.
Scheduled Server Workload	Lists the estimated scheduled workload for the next 24-hour period by server. See “Scheduled Server Workload report” on page 837.
Scratch Media Availability	Lists the aging distribution of media. Shows how many media are available for overwrite and when other media will become available for overwrite. See “Scratch Media Availability report” on page 838.
Test Run Results	Lists the results for the test run jobs that are set for the selected time period and Backup Exec servers. See “Test Run Results report” on page 839.

Table 18-3 Backup Exec standard reports (*continued*)

Report Name	Description
User Role Details	Lists the roles that are assigned to users for role-based access control. See “User Role Details report” on page 839.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Alert History report

The Alert History report lists all the alerts in the Alert History chronologically, displaying the most recent alerts first.

Table 18-4 Alert History report

Item	Description
Time	Date and time the alert occurred.
Received	Time the alert occurred.
Responded	Time when the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	The name of the job that is associated with the alert.
Backup Exec server	Name of the Backup Exec server on which the alert occurred.
Category	Title of the alert, such as Service Start or Job Failed.
Message	A description of the event that caused the alert.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Alert History By Backup Exec Server report

The Alert History by Backup Exec server report lists all alerts in the alert history, grouped and filtered by Backup Exec server, displaying the most recent alerts first.

Table 18-5 Alert History by Backup Exec Server report

Item	Description
Backup Exec server	Name of the Backup Exec server on which the alert occurred.
Time	Date and time the alert occurred.
Received	Time the alert occurred.
Responded	Time when the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	Name of the job that is associated with the alert.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Anomaly Detection Summary report

The Anomaly Detection Summary report lists all anomalies that are detected for a backup job.

By default, 90 days data is displayed in the report.

Table 18-6 Anomaly Detection Summary report

Item	Description
Detection Time	Date and time when the anomaly is detected.
Image Size	Size of the backup image. Usual size range of the image is displayed.
Backup Files Count	Number of items that are backed up. Usual range of the number of items that are backed up is displayed.
Data Transferred	Total size of data transferred (in MB) over the network. Usual size range of the data transferred over the network is displayed.

Table 18-6 Anomaly Detection Summary report (*continued*)

Item	Description
Backup Time	Backup time in seconds. Usual range of the backup time is displayed.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication. Usual range of the deduplication ratio is displayed.
Severity	Severity of the anomaly. ■ Low ■ Medium ■ High
Anomaly Status	Status of the anomaly. ■ Not Reviewed ■ False Positive ■ Anomaly Confirmed
Reason	Reason given for the False Positive anomaly status.

Audit Log report

The Audit Log report lists the contents of the audit logs for the selected servers for the selected time period.

Table 18-7 Audit Log report

Item	Description
Category	Category in which the change occurred, such as Logon Account, Alerts, or Job.
Date Entered	Time and date the change occurred.
Message	Description of the change that was made in Backup Exec.
User Name	User that made the change.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Backup Job Success Rate report

The Backup Job Success Rate report lists the success rate for backup jobs.

Table 18-8 Backup Success Rate report

Item	Description
Server	Name of the server being backed up.
Date	Date the backup job was processed.
Total Jobs	Total number of jobs that were processed by the Backup Exec server.
Successful	Total number of jobs successfully performed by the Backup Exec server.
Success Rate	Percentage of successful jobs that were processed by the Backup Exec server.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Backup Recommendations report

The Backup Recommendations report lists any recommendations that can help you to better manage your backups. The recommendations may include better ways to back up specific types of data or suggestions for how to use other products.

Table 18-9 Backup Recommendations report

Item	Description
Backup Exec server	Name of the Backup Exec server for which the recommendation applies.
Job Name	Name of the job that is associated with the recommendation.
Start Time	Date and time when the job that is associated with the recommendation ran.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Backup Resource Success Rate report

The Backup Resource Success Rate report lists the success rate for backup jobs for a specific number of days for resources on selected servers.

Table 18-10 Backup Resource Success Rate report

Item	Description
Resource	Name of the system being backed up.
Date	Date the backup job was processed.
Total Jobs	Total number of jobs that were processed by the Backup Exec server.
Successful	Total number of jobs successfully performed by the Backup Exec server.
Success Rate	Percentage of successful jobs that were processed by the Backup Exec server.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Backup Sets by Media Set report

The Backup Sets by Media Set report lists all the backup sets by media set.

Table 18-11 Backup Sets by Media Sets report

Item	Description
Media Set	Name of the media set on which the job ran.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Method	Specific type of backup.
Business-Critical	Indicates whether the backup set includes an item that was tagged as a business-critical resource.
Date / Time	Date and time the data was backed up.
Backup Set Description / Source	Describes the data that was backed up and the location of the data.
Directories	Number of directories that were backed up.

Table 18-11 Backup Sets by Media Sets report (*continued*)

Item	Description
Files	Number of files that were backed up.
Size	Amount of data that was backed up.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Backup Size By Resource report

The Backup Size By Resource report lists the backup size for each resource job for up to seven previous jobs. It also computes the trailing average, which is the average of the amount of data that was backed up in the seven previous jobs.

Table 18-12 Backup Size by Resource Job report

Item	Description
Server	Name of the Backup Exec server where the data for the backup job was located.
Resource	Name of the resource that was backed up.
Job	Name of the backup job.
Job Date and Time Run	Date and time the backup job was processed.
Backup Size	Amount of data that was backed up.
Trailing Avg	Average amount of data that was backed up during the seven previous runs.
Difference %	Amount by which the data that was backed up in the current job differs from the data that was backed up in the previous backup jobs.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Cloud Storage Summary report

The Cloud Storage Summary report displays a summary of the size of cloud storage that the backed up data uses on the Backup Exec server.

Table 18-13 Cloud Storage Summary report

Item	Description
Device Name	The name of the cloud storage device.
Cloud Storage Server	The fully qualified name of the server on which the device exists.
Cloud Bucket	The name of the storage location on the cloud storage device. These storage units are called buckets..
Bytes Written	The amount of data written to the cloud storage device.
Bytes Read	The amount of data read from the cloud storage device.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Daily Device Utilization report

The Daily Device Utilization report lists the percentage of the storage devices’ capacity that the Backup Exec server uses.

Table 18-14 Daily Device Utilization report

Item	Description
Drive Name	Name of the storage device and the Backup Exec server where the device is located.
Status	Status of the storage device Statuses for storage devices are as follows: ■ Pause The storage device is temporarily stopped. ■ Enable The storage device is available for use with Backup Exec. If the storage device is disabled, it is available for use with other applications. ■ Online The storage device is available for use. ■ Offline Backup Exec cannot access the storage device.

Table 18-14 Daily Device Utilization report (*continued*)

Item	Description
Date	Date the storage device was used.
Jobs	Number of jobs that were processed by the Backup Exec server's storage device.
Size	The amount of data that was processed by the Backup Exec server's storage device.
Utilization (%)	Percentage of device utilization.

See ["Running a report now"](#) on page 787.

See ["Creating a custom report"](#) on page 789.

Deduplication Disk and Cloud Deduplication Device Summary report

The Deduplication Disk and Cloud Deduplication Device Summary report displays a summary of the deduplication operations for local deduplication disk storage, shared deduplication disk storage, and cloud deduplication disk storage.

Table 18-15 Deduplication Disk and Cloud Deduplication Device Summary report

Item	Description
State	Device state, such as online and enabled.
Created	Date media was created.
Total Capacity	Total capacity of the deduplication disk storage.
Used Capacity	Capacity presently used by the deduplication disk storage.
Available Capacity	Remaining capacity of the deduplication disk storage.
Percent Full	Percentage of storage space that is available in the deduplication disk storage.
Protected Bytes	Total amount of data that is selected for backup in all jobs using the device before deduplication occurs.

Table 18-15 Deduplication Disk and Cloud Deduplication Device Summary report (*continued*)

Item	Description
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Deduplication Summary report

The Deduplication Summary report displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec server.

Table 18-16 Deduplication Summary report

Item	Description
Job Name	Name of the job.
Start Time	Time of day that Backup Exec attempted to start the job.
Duration	Length of time the operation took to process.
Size	The amount of data processed.
Size/Minute	Number of kilobytes, megabytes, or gigabytes processed per minute.
Scanned Byte Count	Total amount of data that is selected for backup before deduplication occurs.
Stored Byte Count	The amount of unique data is stored after deduplication occurs.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.
Status	Status of the operation, such as Completed (Success), Failed, or Canceled.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Device Summary report

The Device Summary report lists all the devices for each selected Backup Exec server.

Table 18-17 Device Summary report

Item	Description
Server	Name of the server where the device is located.
Drive Name	Name of the drive in the robotic library.
Target	Address of the storage device that is connected to the Backup Exec server.
State	Device state, such as online.
Created	Date on which the media was created.
Cleaned	Date on which the last cleaning job was run on the drive.
Hours	Hours the device has been in use since the last cleaning job.
Errors	Number of errors occurring since the last cleaning job.
Size	Amount of data that was read and written since the last cleaning job.
Mounts	Number of mounts occurring since the last cleaning job.
Hours	Total number of hours the device has been in use.
Errors	Total number of errors occurring on the device.
Size	Amount of data that was read and written to the device.
Mounts	Total number of mounts occurring to the device.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Disk Storage Summary report

The Disk Storage Summary report displays disk usage statistics for Backup Exec server disk storage.

Table 18-18 Disk storage summary report

Item	Description
Device Name	Name of the disk storage device.
State	State of the device. Device states include the following: ■ Online ■ Enabled ■ Offline ■ Paused ■ Disabled
Local Access Path	Path on the disk where backup data is stored.
Total Capacity	Total capacity of the disk.
Used Space	Amount of disk space being used as storage.
Free Space	Amount of disk space remaining.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Error-Handling Rules report

The Error-Handling Rules report lists all error-handling rules and provides details about each rule.

Table 18-19 Error Handling Rules report

Item	Description
Rule Name	Name of the Error-Handling rule.
Notes	Information that was entered in the Notes section when the error-handling rule was created.
Job Status	Final job status that activates the rule. Possible statuses are as follows: ■ Error ■ Canceled

Table 18-19 Error Handling Rules report (*continued*)

Item	Description
Error Category	Category of error for which the rule will be applied. Available error categories include the following: ■ Device ■ Job ■ Media ■ Network ■ Other ■ Resource ■ Security ■ Server ■ System
Enabled	Displays if the rule is enabled or disabled.
Cancel Job	Displays an X if this option is selected for the error-handling rule. The option cancels all jobs after the maximum number of retries have been attempted.
Pause Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to pause the job until you can manually clear the error.
Retry Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to retry the job.
Maximum Retries	Number of times the job is to be retried.
Retry Interval (minutes)	Number of minutes Backup Exec waits before trying to run the job again.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Event Recipients report

The Event Recipient report lists events that were registered for each notification recipient.

Table 18-20 Event Recipients report

Item	Description
Recipient Name	Name of the recipient.
Recipient type	Designates to whom the Event Recipients report is sent, such as an individual recipient or a group of recipients.
Event Type	Alert category or ad hoc job.
Event Name	Detail for the alert category or ad hoc job.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Failed Backup Jobs report

The Failed Backup Jobs report lists all the failed backup jobs. The jobs are sorted by the server and specified time frame.

Table 18-21 Failed Jobs report

Item	Description
Resource	Name of the system being backed up.
Start Time	Date and time the backup job started.
Duration	Length of time the operation took to process.
Job Name	Name of job that failed.
Business-Critical	Indicates whether the job included items that were tagged as business-critical resources.
Category	Category for the failed job that may be generated by a system, job, media, or device error.
Error Code	Displays the error code that corresponds to the failure.
Description	Describes the event that caused the error.
Status	Status of the operation, such as Error.
Device Name	Name of the device on which the job ran.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Jobs Summary report

The Jobs Summary report lists all jobs that have run within the specified time range. The jobs are listed in chronological order.

Table 18-22 Jobs Summary report

Item	Description
Start Time	Date and time the operation started.
Job Name	Name of the completed job.
Server	Number of servers that are protected.
Duration	Length of time that the operation processed.
Size	The amount of data processed.
Files	Number of files processed.
Directories	Number of directories processed.
Entities	Number of Microsoft 365 workloads that are protected.
Size/Minute	Number of kilobytes, megabytes, or gigabytes processed per minute.
Skipped	Number of files that were skipped during the operation.
Corrupt Files	Number of corrupt files that were encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed (Success), Failed, or Canceled.
Type	Lists the type of job that Backup Exec ran within the specified time range.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Malware Scan Summary report

The Malware Scan Summary report lists the information about the malware scan of the backup sets that is run as part of the restore job or ad hoc scan job.

By default, data for 90 days is displayed in the report.

Table 18-23 Malware Scan Summary report

Item	Description
Backup Set	Name of the backup set, which includes the backup set time and the type of backup set.
Job Type	Type of the job. In this case, it is Malware Scan for an ad hoc scan job and Restore for a scan that is performed as part of the restore job.
Scan Time	Date and time when the malware scan was run.
Scan Duration	Amount of time that it took to run the malware scan.
Scan Status	Status of the malware scan. ■ Malware Found ■ Malware not Found ■ Scan Incomplete ■ Scan Failed
Malware Count	Number of malware that is detected.
Malware Types	Types of malware that are detected.
Number of Malware Types	Number of malware types that are detected.
Scanned by	The scan engine that was used to perform malware scan.

Managed Backup Exec Servers report

The Managed Backup Exec Servers report lists status and configuration information for the managed Backup Exec servers in a Central Admin Server feature environment.

Table 18-24 Managed Backup Exec Servers report

Item	Description
CAS Servers	Names of the central administration servers.
Managed Backup Exec server	Names of the managed Backup Exec servers.

Table 18-24 Managed Backup Exec Servers report (*continued*)

Item	Description
Status	Status of the server. Possible statuses include the following: ■ Online - available for use. ■ Stalled - not responding immediately to messages ■ No Comm - communications to the server have been lost for some period of time.
Stalled	Time limit that was used for determining Stalled communications status.
No Comm	Time limit that was used for determining No Comm communications status.
Catalog Location	Location of the catalog information. Possible locations are as follows: ■ Local - the catalog information is located on the managed Backup Exec server. ■ CAS - the catalog information is located on the central administration server.
Logs	When job logs are uploaded from the managed server to the CAS database. Possible upload times are as follows: ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never
History	When job history is uploaded from the managed server to the CAS database. Possible upload times are as follows: ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never

Table 18-24 Managed Backup Exec Servers report (*continued*)

Item	Description
Status	When status is uploaded from the managed server to the CAS database. Possible upload times are as follows: ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never
Display Alert	Displays Yes if you have configured an alert to be set if the time difference between the central administration server's clock and a managed Backup Exec server's clock exceeds a preset value (maximum time difference tolerance).
Sec	Maximum time difference tolerance in seconds that is set for server.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Media Audit report

The Media Audit report lists the recent configuration changes that you made to your media.

In a Central Admin Server feature (CAS) environment, if you run this report from the central administration server the report provides data only for the media for the central administration server; it does not provide any data for any of the managed Backup Exec servers. To obtain the media audit data for a managed Backup Exec server, you can do one of the following:

- Log on locally to the managed Backup Exec server and run the report from it.
- Use the Remote Administrator feature to log on to the managed Backup Exec server from a remote Windows server or workstation, and then run the report.

Table 18-25 Media Audit report

Item	Description
Date Entered	Time and date the change occurred.
Message	Description of the change that was made to the media.

Table 18-25 Media Audit report (*continued*)

Item	Description
User Name	User that made the change.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Media Errors report

The Media Errors report lists the number of errors that occur on all media.

Table 18-26 Media Errors report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Business-Critical	Indicates whether the media contains a backup set that has a business-critical resource.
Total Mounts	Total number of times this media has been mounted.
Total In Use Hours	Total number of hours that this media has been in use.
Total Errors	Total number of error alerts for the system, jobs, media, and devices.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Media Required for Recovery report

The Media Required for Recovery report lists the media that contain the backup sets for each system that was backed during the specified time period. However, this report may be inaccurate if media overwrite settings allow the media to be overwritten.

Table 18-27 Media Required for Recovery report

Item	Description
Date	Date and time the backup job set was created.

Table 18-27 Media Required for Recovery report (*continued*)

Item	Description
Media Location Name	Name of the storage device where the media that was used for the backup job is stored.
Media Label	Media label that is assigned to the media.
Business-Critical	Indicates whether the media includes a backup set that has a business-critical resource.
Recycle Time	Displays the date and the time after which the media can be overwritten.
Backup Method	Specific type of backup.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Media Summary report

The Media Summary report lists all media sets and media that are used by Backup Exec servers. Usage statistics are given for each piece of media.

Table 18-28 Media Summary report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Media Type	Type of media cartridge, such as 4mm.
Allocated	Date media was allocated to a media set as a result of an overwrite operation.
Modified	Date data was last written to the media.
Business-Critical	Indicates whether the media includes a backup set that has a business-critical resource.
Hours	Total number of hours that the media has been in use.
Mounts	Total number of times the media has been mounted.
Soft Errors	Number of recoverable read errors encountered.
Hard Errors	Number of unrecoverable read errors encountered.

Table 18-28 Media Summary report (*continued*)

Item	Description
Write Size	Amount of data that has been written to the media.
Current Size	Estimate of the amount of data currently on the media.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Media Vault Contents report

The Media Vault Contents report lists all the media in a specified media vault.

Table 18-29 Media Vault Contents report

Item	Description
Vault Name	Location of the media.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Media Set	Name of media set to which the media belongs.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Move Media to Vault report

Lists all media that you can move to a media vault.

The media listed are not currently in a media vault and meet one of the following criteria:

- The media has met or exceeded the vault move date that was specified for the media containing the media.

- The append period has expired, but the overwrite protection period is still current (allocated).

Table 18-30 Move Media to Vault report

Item	Description
Backup Exec server	Name of the Backup Exec server where the data for the backup job was located.
Media Set	Name of the media set.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Location	Location of the media.
Append Period End Date	Last date that data may be added to the media.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Vault Name	Name of vault to which media is to be moved.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Operations Overview report

The Operations Overview report lists details for past and future Backup Exec operations.

Table 18-31 Operations Overview report

Item	Description
Job summary for jobs completed in the past x Hours	Details Backup Exec job activity for the specified time period.
Errors	Total number of error alerts for the system, jobs, media, and devices.

Table 18-31 Operations Overview report (*continued*)

Item	Description
Warnings	Total number of warning alerts for of jobs, media, and devices.
Information	Total number of informational alerts for the system, jobs, media, and devices.
Attention Required	Total number of alerts that require a response from the user.
Completed (Failed)	Total number of jobs that failed.
Completed (Canceled)	Total number of canceled jobs.
Completed (Success)	Total number of jobs that completed successfully.
Exceptions	Total number of jobs that completed successfully, but may contain one or more skipped files, corrupt files, virus infected files or files in use.
Total Data Backed Up	Total amount of data that was backed up in kilobytes, megabytes, or gigabytes.
Total Media Used	Total number of media that were used to back up the completed jobs.
Missed	Total number of missed jobs.
Recovered	Total number of recovered jobs.
Active Jobs	Total number of active jobs.
Scheduled Jobs	Displays those jobs whose scheduled start times begin within 72 hours of the job being created. Jobs with recurring schedules also appear if their start times begin within 72 hours of the job's last start time.
Jobs On Hold	Total number of jobs on hold.
Job Status	The status of the jobs.
Scratch Media	Total number of scratch media available.
Recyclable	Total number of recyclable media available.
Imported	Number of imported media. Imported media is media that was created by a product other than this installation of Backup Exec.
Allocated	Number of allocated media (media belonging to a user media set).

Table 18-31 Operations Overview report (*continued*)

Item	Description
Total Overwritable Media	Total number of overwritable media available.
Total Appendable Media	Total number of appendable media available.
Media Overwrite Protection Level	Displays the level of overwrite protection (Full, Partial, None) that is assigned to the media.
Online Devices	Total number of online devices.
Offline Devices	Total number of offline devices.
Disabled Devices	Total number of disabled devices.
Paused Devices	Total number of paused devices.
Disabled	Lists the name of the devices that are disabled.
Paused	Name of the paused devices.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Overnight Summary report

The Overnight Summary report lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were due to run but did not run. Jobs are given a grace period of 24 hours before being marked as past due.

Table 18-32 Overnight Summary report

Item	Description
Resource	System being backed up.
Type	Displays the type of job that Backup Exec runs to produce the Overnight Summary report. Because the Overnight Summary report lists the results of backup jobs for each resource during the past 24 hours, Backup is always the type of job that appears

Table 18-32 Overnight Summary report (*continued*)

Item	Description
Start time	Date and time the operation started.
Business-Critical	Indicates if the item was tagged as a business-critical resource.
Status	Status of the operation.
Error Category	Category for the job that may be generated by a system, job, media, or device error.
Backup Exec server	Name of the Backup Exec server on which the job ran.
Device Name	Name of the device on which the job ran.
Total Tasks	Total number of jobs that ran within the last 24 hours.
Uncorrected Exceptions	Number of the jobs that fail and were not run again with successful completion.
Service Level	Percentage of jobs that ran successfully.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Problem Files report

The Problem Files report lists all the problem files that were reported for jobs. The files are grouped by day and resource.

Table 18-33 Problem Files report

Item	Description
Date	Date the problem file was encountered.
Resource	System on which the problem file is located.
Time	Time the problem file was encountered.
Reason	Error code that is listed in the job log summary.
File Name	Name of the problem file.
Type	Lists the type of job that Backup Exec ran when problematic files were detected.

Table 18-33 Problem Files report (*continued*)

Item	Description
Backup Exec Server	Name of the Backup Exec server on which the file is located.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Recently Written Media report

The Recently Written Media report lists all the media that has been modified within the specified period.

Table 18-34 Recently Written Media report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Location	Location of the media, such as the storage vault name or drive name.
Set	Name of backup set.
Date and Time Modified	Date and time media was last modified.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Recovery Ready Validation Summary

The Recovery Ready Validation Summary report shows the summary of the virtual machines that you validated.

Table 18-35 Recovery Ready Validation Summary report

Item	Description
Backup Set Details	Displays the details of the backup set. name of the backup set, size of the backup set, and name of the device ■ Backup set: Name of the backup set. ■ Size: Size of the backup set. ■ Device Name: Name of the device.
Job Statistics	Displays details about the validation job. ■ Job Name: Name of the validation job. ■ Start Time: Time when the validation job started.
Backup Set Validation	Displays the status of the tests and checks run on the backup set. ■ Register VM: Status of the Register VM test. ■ Power ON: Status of the Power ON test. ■ Heartbeat Check: Status of the Heartbeat check.
Validation Status	Displays the final validation status of the backup set. ■ Validation passed ■ Validation failed ■ Unable to validate

The validation status and the range of the report that you selected are displayed in the Recovery Ready Validation Summary report.

If you do not select the range, the report is displayed for last 30 days. If you run the same validation job multiple times for the same backup set, the report displays information for the latest run of the validation job.

Resource Protected Recently report

The Resource Protected Recently report lists all job detail statistics and exceptions that occurred on a Backup Exec server for which you run this report.

Table 18-36 Resource Protected Recently report

Item	Description
Start Time	Date and time the backup job started.
Business-Critical	Indicates whether the item was tagged as a business-critical resource.

Table 18-36 Resource Protected Recently report (*continued*)

Item	Description
Duration	Amount of time that was required for the job to complete.
Size	Amount of data that was backed up.
Files	Number of files that were backed up.
Directories	Number of directories that were backed up.
Size per Min	Amount of data that was backed up per minute.
Skipped	Number of files that were skipped during the backup.
Corrupt Files	Number of corrupt files that were detected during the backup.
Files in Use	Number of files that were in use during the backup.
Status	Status of the backup job.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Resource Risk Assessment report

The Resource Risk Assessment report shows job information for resources on which the last backup job that was run on the resource failed.

Table 18-37 Resource Risk Assessment report

Item	Description
Resource	System on which the job ran.
Error Text	Describes the event that caused the job to fail.
Start Time	Time the operation started.
Job	Name of the job that failed.
Business-Critical	Indicates whether the item was tagged as a business-critical resource.

Table 18-37 Resource Risk Assessment report (*continued*)

Item	Description
Error Category	The category for the failed job that may be generated by a system, job, media, or device error.
Backup Exec Server	Name of the Backup Exec server on which the job ran.
Device Name	Name of the device on which the job ran.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Restore Set Details by Resource report

The Restore Set Details by Resource report lists all restore jobs that ran within the specified time range on a selected server. The jobs are grouped by the server and resource.

Table 18-38 Resource Set Details by Resource report

Item	Description
Resource	Name of the system being backed up.
Start Time	Date and time the operation started.
Duration	Length of time the operation took to process.
Size	The amount of data processed.
Files	Number of files processed.
Directories	Number of directories processed.
Data/Minute	Amount of data that was processed per minute.
Skipped	Number of files that were skipped during the operation.
Corrupt Files	Number of corrupt files that were encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Retrieve Media from Vault report

The Retrieve Media from Vault report lists all reusable media currently in a specified media vault.

Table 18-39 Retrieve Media from Vault report

Item	Description
Cartridge Label	Displays the name of the disk cartridge. Disk cartridge names cannot exceed 128 characters. You can rename the disk cartridge. See “Editing disk cartridge properties” on page 346.
Vault Name	Displays the name of the vault where the media is located.
Media Set Name	Displays the name of the media set.
Offsite Return Date	Displays the date that the media was returned to the off-site vault.
Recycle Date	Displays the date after which the media can be overwritten.
No Append Date	Displays the date on which Backup Exec can no longer append data to the media.
Rule Name	Displays the name of vault media rule that is applied to the media.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Robotic Library Inventory report

The Robotic Library Inventory report lists the contents of slots in robotic libraries that are attached to Backup Exec servers. Usage statistics are provided for each piece of media.

Table 18-40 Robotic Library Inventory report

Item	Description
Server	Name of the server where the robotic library is located.
Device Name	Name of the robotic library.
Slot	Sequential number of the slot in the robotic library.

Table 18-40 Robotic Library Inventory report (*continued*)

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
State	State of operation of the slot: paused, disabled, enabled, offline, or online.
Modified	Date the media in the slot was last accessed.
Business-Critical	Indicates whether the media includes a backup set that has a business-critical resource.
Write	Number of bytes that have been written to this media.
Full	Space available on a media; "1" indicates that media is full and "0" indicates that there is space available on the media.
Hours	Total number of hours this media has been in use.
Mounts	Total number of times this media has been mounted.
Append	The time remaining in the media's append period.

See ["Running a report now"](#) on page 787.

See ["Creating a custom report"](#) on page 789.

Scheduled Server Workload report

The Scheduled Server Workload report displays the estimated scheduled workload for a server during the next 24-hour period or a user-defined time period. The report only displays recurring jobs that have already run at least one time, not jobs scheduled to run once.

Table 18-41 Scheduled Server Workload report

Item	Description
Backup Exec server	Name of the Backup Exec server that will process the scheduled jobs.
Job	Name of the job that is scheduled to run.
Next Due Date	Time and day the next job is scheduled to run.
Backup Size	Estimated amount of data to be processed during the next 24 hours.

Table 18-41 Scheduled Server Workload report (*continued*)

Item	Description
Total Size	Total amount of data to be processed on the server during the next 24 hours.
Total Size	Total amount of data to be processed on all Backup Exec servers.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Scratch Media Availability report

The Scratch Media Availability report shows the aging distribution of media, how many media are available for overwrite, and when other media will become available for overwrite.

Table 18-42 Scratch Media Availability report

Item	Description
Cartridge Label	Cartridge label that was assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label You can rename the cartridge. See “Editing disk cartridge properties” on page 346.
Media Location Name	Name of the storage device that contains the actual media.
Total Capacity	Total native capacity of the scratch media without using compression.
Append Hours Remaining	Capacity of scratch media available for append.
Remaining Capacity	Total amount of remaining native capacity of the scratch media without compression.
Retention Hours Remaining	The amount of time remaining to retain and protect the media from being overwritten.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

Test Run Results report

The Test Run Results report displays the results for the test run jobs that are set for the selected period.

Table 18-43 Test Run Results report

Item	Description
Backup Exec server	Name of the Backup Exec server on which the job ran.
Job Date and Time Run	Date and time the backup job was processed.
Job Name	Name of the test run job.
Backup Sets	Name of the backup set.
Credential Check	Indicates if the Backup Exec logon account was verified as correct for the resources being backed up.
Backup Size	Size in kilobytes, megabytes, or gigabytes of the backup.
Media Type	Type of media used, such as 4mm.
Device Name	Name of the device, such as the name of the robotic library.
Max Needed	Amount of space that is needed on the media to run the job.
Online	Capacity of media available in the device to which data can be appended.
Media Total	Total amount of appendable media available to the system.
Online	Capacity of media available in the device to which data can be written.
Media Total	Total amount of overwritable media available to the system.

See [“Running a report now”](#) on page 787.

See [“Creating a custom report”](#) on page 789.

User Role Details report

The User Role Details report lists all the users that are added to the role-based access control list and the roles that are assigned to each user.

Table 18-44 User Role Details report

Item	Description
User Name	Name of the user in the role-based access control list.
Roles	Role that is assigned to each user. The following roles can be assigned to a user: ■ Administrator ■ Storage Administrator ■ Backup Administrator ■ Restore Administrator ■ View Only
Synchronized	Status of the user who is synchronized from CAS to MBES in a CASO environment. In a non CASO environment, the status is displayed as Not Applicable.

GDPR Guard

This chapter includes the following topics:

- [About GDPR Guard](#)
- [Backup Exec Management Command Line \(BEMCLI\) commands for import and export](#)
- [Supported types of backed up data](#)
- [How to block access to backed up items](#)
- [Restoring blocked items](#)
- [Best practices for blocking access to backed up items with GDPR Guard](#)

About GDPR Guard

An organization may require to block some of the backed-up items for privacy and compliance reasons, such as General Data Protection Regulation (GDPR). These items should not be viewed or restored.

Backup Exec provides a way to import the list of items to block. You can use the GDPR Guard feature to specify the list of backed-up items whose access needs to be blocked.

An organization can use any tool to create a list of the blocked items. Backup Exec take this list in the form of a CSV file. The format of the CSV file used for specifying blocked items information during import is generic and can accommodate any CSV file generated by different tools.

To import the CSV file containing blocked items into Backup Exec, you must use the `Import-BEItemsToBlock` BEMCLI command. You can import CSV files multiple times into Backup Exec and each CSV can contain blocked items from multiple servers.

You can export the content of all imported files into one CSV file using the `Export-BEItemsBlocked` BEMCLI command.

After the blocked items are imported, the restore browse and search view do not display the blocked items. When you run a restore job, the blocked items are not available for restore. The blocked items continue to be part of backups and are not deleted from the backup media.

Backup Exec ensures that information about blocked items is protected using encryption and all operations that are related to blocked items are recorded in the audit log for compliance requirements.

When you run a restore job, you can use the **Allow restore of blocked items** option to restore blocked items. When you select this option, you must provide a reason to restore the blocked items and that is recorded in the audit log. Backup Exec allows only an SLA owner to restore the blocked items.

Following are the key features of GDPR Guard:

- Blocked items cannot be viewed or restored from the Backup Exec console.
- Backup Exec ensures integrity and protection of the blocked items data.
- The file format that is used for blocked files operations (import and export) is CSV, which is an accepted and easy to use format. A CSV file supports all types of character encoding.
- All operations that are related to blocked items are recorded in audit logs and Windows event log, which can be used for compliance requirements.
- The imported blocked items on Central Administration Server (CAS) are auto-synced across all Managed Backup Exec Server (MBES)s that support blocking of items, ensuring blocking in the CAS-MBES environment without actually importing blocked items in each MBES.

Note: The import command cannot be run on MBES. The export command can be run on CAS and MBES.

Figure 19-1 Workflow for blocking access to backed up items with GDPR Guard



See [“ Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

See [“Supported types of backed up data”](#) on page 845.

See [“How to block access to backed up items”](#) on page 846.

See [“Restoring blocked items”](#) on page 847.

See [“Best practices for blocking access to backed up items with GDPR Guard”](#) on page 848.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

Backup Exec Management Command Line (BEMCLI) commands for import and export

Blocked items can be specified and viewed using the following BEMCLI commands.

- Import-BEItemsToBlock
- Export-BEItemsBlocked

For more information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:<Backup Exec install path>\Backup Exec

Import blocked items

The `Import-BEItemsToBlock` BEMCLI command is used to provide list of blocked items to Backup Exec. This command requires full path of the CSV file as one of its parameters.

The CSV file contains information about the items to be blocked, specifically the server name and the full path. You can use the import command to add new entries to the existing list of blocked items used by Backup Exec.

The column names in the CSV file are not required to be in a specific order and the import command can accept a CSV file, which can contain extra columns. This makes the format of the CSV file generic to accommodate any CSV file generated by different tools.

In the same operation you can import the blocked items that belong to different servers and are part of the same CSV.

Every time the import command is run, an entry is logged into the audit log under the **Compliance** category and the Windows Event Viewer.

Export blocked items

You can use the `Export-BEItemsBlocked` BEMCLI command to export information about all the blocked items that are imported.

Information about the blocked items is exported to a CSV file at a specified location. This export command requires the location to create the CSV file, as a parameter. The export command creates a CSV file in the same format as the input CSV file.

Every time the export command is run, an entry is logged into the audit log under the **Compliance** category and the Windows Event Viewer.

It is recommended to regularly back up the imported blocked items by running the export BEMCLI command.

CAS-MBES behavior for import and export commands

Import command behavior:

- Import command can be run only on CAS and is blocked on MBES.
- After the import operation on CAS is complete, the blocked items information along with the encryption key is sent to MBES.

- If an MBES is offline during the import operation, the blocked items' information is shared with that MBES after it is online.
- In case of a rolling upgrade, older MBES does not receive the blocked items. After MBES is upgraded to the current version of Backup Exec, the blocked items information is synced with MBES.
- If a new MBES is added, the encryption key and blocked items information is synced with the MBES.
- If there is an update to the blocked items or reset operations, all information including CSV data is synced with MBES.

The export command runs on both CAS and MBES.

See [“About GDPR Guard”](#) on page 841.

See [“Supported types of backed up data”](#) on page 845.

See [“How to block access to backed up items”](#) on page 846.

See [“Restoring blocked items”](#) on page 847.

See [“Best practices for blocking access to backed up items with GDPR Guard”](#) on page 848.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

Supported types of backed up data

Backup Exec supports the following types of backed up data:

- File systems that Backup Exec supports. For example, NTFS, Linux.

Note: NDMP file servers are not supported.

- Windows shares
- File and folder data of virtual machine backups; both VMware and Hyper-V.

See [“About GDPR Guard”](#) on page 841.

See [“Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

See [“How to block access to backed up items”](#) on page 846.

See [“Restoring blocked items”](#) on page 847.

See [“Best practices for blocking access to backed up items with GDPR Guard”](#) on page 848.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

How to block access to backed up items

You can block access to backed-up items from Backup Exec. Backup Exec takes as input the list of items that need to be blocked from restore. Backup Exec honors the blocked items during restore and ensures integrity and protection of the blocked items' information.

Ensure that you have a list of the items that need to be blocked from restore. You can use any tool to create a list of the blocked items. The list must be in a CSV file format and contain the item path and server information for every blocked item.

All import and export command-related operations and restore job executions are recorded in the audit log and Windows Event Log.

The restore job log contains the information that the blocked items are not restored. The actual file names are not listed.

To block access to backed up items

- 1 Import the CSV file containing the list of blocked items in Backup Exec using the `Import-BEItemsToBlock` BEMCLI command.

You can export this imported list of blocked items to a different location by using the `Export-BEItemsBlocked` BEMCLI command. The exported CSV file is your backup of blocked items.

See [“Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

- 2 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 3 Select **Files, folders, or volumes**, and then click **Next**.

- 4 Select **File and folder backups to a point-in-time, File and folder backups from a backup set, or Files and folders located through Search**.

The list of blocked items that were imported are not available when you select files and folders. When the blocked items have to be restored, Backup Exec allows only the SLA owner to restore these items and reason for restore is recorded in the audit log.

- 5 Follow the Restore Wizard prompts to restore the data.

See [“Restoring file system data”](#) on page 244.

See [“About GDPR Guard”](#) on page 841.

See [“Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

See [“Supported types of backed up data”](#) on page 845.

See [“Restoring blocked items”](#) on page 847.

See [“Best practices for blocking access to backed up items with GDPR Guard”](#) on page 848.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

Restoring blocked items

Blocked items, by default are not displayed in restore and search view and cannot be restored.

In scenarios, when the blocked items have to be restored, Backup Exec permits only the SLA owner to restore these items and reason for restore is recorded in the audit log. The audit log records the reason for which the blocked files are restored.

If you want to know the owner of the system logon account, click the Backup Exec button, and then select **Configuration and Settings > Select Logon Accounts > Manage Logon Accounts**. On the **Logon Account Management** dialog box, the **Owner** column displays the owner of the system logon account.

To restore the blocked files

- 1 Log on as an SLA owner.
- 2 In the Restore Wizard, select **Files, folders, or volumes**, and then click **Next**.
- 3 Select the **Allow restore of blocked items** check box.
- 4 In **Restore Blocked Items** dialog box, type the reason for restore.

5 In the **Restore Wizard**, click **Next** to view the blocked files in the **Restore** view. The restore view lists the blocked items that can be selected for restore.

6 Follow the Restore Wizard prompts to restore the data.

See [“About GDPR Guard”](#) on page 841.

See [“Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

See [“How to block access to backed up items”](#) on page 846.

See [“Supported types of backed up data”](#) on page 845.

See [“Best practices for blocking access to backed up items with GDPR Guard”](#) on page 848.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

Best practices for blocking access to backed up items with GDPR Guard

Review the following best practices before you block access to backed up items:

- While specifying the blocked item in the CSV file for import, ensure that you use '*' related wildcards to specify a set of items within a folder. For example, H:\Folder1*.txt and E:*. * can be used for blocked items path in the CSV file to block all text files under H:\Folder1 or to block all files in the E: directory.
- For the blocked items that belong to a non-Windows computer, the file path is case-sensitive. Ensure that you provide a file path with the correct case.
- Ensure that the server names mentioned for blocked item entry in the CSV file match with the server name appearing in the **Servers** list on the Media server (**Backup and restore Tab > Servers Column**). If the fully qualified domain (FQDN) server name is used for a blocked item entry in the CSV file and the server appears with a different name in the **Servers** list, ensure that the remote agent of that server publishes itself on the media server. If the server name does not match, entries are skipped.
- Run the export BEMCLI command after the import operation is complete. The exported file is your backup for all the blocked items that are imported.
- Enter the reason for restoring blocked items when you create a restore job with the **Allow restore of blocked items** option. The reason is recorded in the audit logs and appears in the audit reports. The reports can be referred for compliance and audit purposes.

- Run the import command again when a media is moved to a new media server and this media has backup sets that contain blocked items. If you do not run the import command again, the blocked items will be available for restore.
- In a CASO environment, run CAS and all MBES with Backup Exec 20.3 or later version for items to be blocked across all media servers. When you run the import command on the CAS server, the information is automatically synced on the MBES server.
- Run the audit log report regularly to keep a record of all the operations related to blocked items.
- Ensure that the CSV used for specifying blocked items during import operation is with one of the following character encodings:
 - Locale encoding which corresponds to the ANSI and the OEM code pages.
 - UTF-8 with Byte-Order-Mark (BOM).
 - UTF-16 Little-Endian with BOM.
 - UTF-16 Big-Endian with BOM.

The import command does not function properly if the input CSV file contains strings from multiple locales.
- If you are restoring a virtual machine or doing a local restore of your computer using Simplified Disaster Recovery, ensure that you delete the blocked items manually or using a post-processing script once the restore is complete. Otherwise, blocked items are restored in such cases.

See [“About GDPR Guard”](#) on page 841.

See [“Backup Exec Management Command Line \(BEMCLI\) commands for import and export”](#) on page 843.

See [“Supported types of backed up data”](#) on page 845.

See [“How to block access to backed up items”](#) on page 846.

See [“Restoring blocked items”](#) on page 847.

See [“Troubleshooting blocked access to backed up items with GDPR Guard”](#) on page 865.

Troubleshooting Backup Exec

This chapter includes the following topics:

- [Troubleshooting hardware-related issues in Backup Exec](#)
- [Troubleshooting robotic libraries and tape drives](#)
- [How to get more information about alerts and error messages](#)
- [Troubleshooting backup issues in Backup Exec](#)
- [Troubleshooting failed components in the SAN](#)
- [Troubleshooting installation issues in Backup Exec](#)
- [Troubleshooting blocked access to backed up items with GDPR Guard](#)
- [How to improve Backup Exec's performance](#)
- [Accessing Arctera Online](#)
- [Searching the Arctera Knowledge Base](#)
- [Contacting Backup Exec Technical Support](#)
- [Using Remote Assistance](#)
- [Managing your Backup Exec support cases](#)
- [About Backup Exec diagnostic tools](#)
- [Running the Veritas QuickAssist Help Tool](#)
- [Generating a diagnostic file for troubleshooting Backup Exec](#)

- [Running the begather utility to troubleshoot Backup Exec components on Linux servers](#)
- [Using the Backup Exec Debug Monitor for troubleshooting](#)
- [About the Backup Exec debug tool](#)

Troubleshooting hardware-related issues in Backup Exec

For common hardware-related issues, review the following frequently asked questions:

Table 20-1 Hardware-related questions

Question	Answer
How do I know if my storage device is supported?	You can find a list of compatible devices at the Backup Exec Hardware Compatibility List. If your drive is listed on the hardware compatibility list, run the Configure Storage wizard and install device drivers. From the Configure Storage wizard, the Device Driver Installation Wizard finds and installs the most suitable driver for your tape drive.
How can I troubleshoot issues with a robotic library or tape drive?	The following sections provide tips for troubleshooting and configuring tape device and robotic library hardware: See “Troubleshooting robotic libraries and tape drives” on page 853. See “Starting and stopping Backup Exec services” on page 771. See “Deleting a storage device” on page 558.

Table 20-1 Hardware-related questions (*continued*)

Question	Answer
I'm getting an error "Storage device [device] reported an error on a request to read or write data to or from media. Error reported: Data error (cyclic redundancy check)." What should I do?	Many factors can cause the cyclic redundancy check (CRC) error. The following list contains the most common reasons for this error and potential ways to resolve the problem: ■ Contaminated read and write heads of the tape device. Check with the hardware manufacturer for proper cleaning techniques. ■ Bad media. Replace the media. Try a new tape that the hardware manufacturer has certified. ■ Tape driver. Load the appropriate Backup Exec tape driver. You can find a list of compatible devices in the Backup Exec Hardware Compatibility List. ■ Wide negotiation for the SCSI controller is not configured properly. If the device is a wide (68 pin) SCSI device, then wide negotiation may and should be used. If the device is a narrow (50 pin) SCSI device, disable wide negotiation. Use the manufacturer's SCSI installation program to disable wide negotiation on the SCSI controller card. ■ SCSI controller transfer rate is too fast. Use the manufacturer's SCSI installation program to lower the SCSI transfer rate. Check with the manufacturer of the controller and the device for the proper configuration for the SCSI transfer rate. ■ SCSI controller synchronous negotiation enabled. Use the manufacturer's SCSI installation program to disable synchronous negotiation on the SCSI controller card. Check with the manufacturer of the controller and the device for the proper configuration for SCSI synchronous negotiation. ■ Incorrect termination or bad cables. Verify that the SCSI cable is good and that it is configured to provide proper SCSI termination. Do not mix passive and active termination. ■ Confirm that the tape drive functions properly. Check with the tape drive manufacturer for diagnostic software to test the condition of the tape drive hardware. ■ General SCSI problems. Isolate the tape drive on its own controller card or try a different SCSI card.

Table 20-1 Hardware-related questions (*continued*)

Question	Answer
Why does my DLT tape drive pause when it catalogs some tapes?	The DLT tape drive maintains internal information about the tape on a tape directory track. The directory track is updated before the tape is ejected from the drive. If the drive is turned off without ejecting the tape first, this information is lost. Re-generating the tape directory information takes several hours to complete, which makes it seem like the drive is hung. Allow sufficient time for the operation to complete and then eject the tape. Normal operation resumes after the directory track is updated.
A backup to my DLT tape drive is stuck at 99% complete. What should I do?	The backup most likely fails to complete because the storage option Eject the media after job completes is selected, and the tape drive does not support the operation. Some tape drives require you to manually remove the tape, such as Digital Linear Tape (DLT), Linear Tape-Open (LTO), Travan, and Onstream drives. To remedy this situation, either uncheck the storage option Eject the media after job completes, or configure an automatic response to the active alert. See “Configuring network options for backup jobs” on page 203.

Troubleshooting robotic libraries and tape drives

This section contains troubleshooting strategies that can help resolve issues with robotic libraries and tape drives. For best results, perform these in order.

Several of these troubleshooting strategies use the Veritas QuickAssist (VQA) tool.

First, verify that the operating system detects the device. Repeat this process whenever Backup Exec loses the ability to manage external tape hardware. If your tape hardware is in a storage area network (SAN), you may need to cycle the SAN power switch to recognize the device. We recommend plugging the tape hardware into a backup power supply that provides surge protection.

Ensure that you have installed the latest firmware for the tape device, robotic library, and host bus adapter. Contact the hardware manufacturer for the latest firmware for the tape drive or robotic library, and for the latest firmware and device drivers for the host bus adapter.

Verify that your license and your installed features support the number of drives you need to use. Depending on your Backup Exec edition and license, you may need to purchase support for additional tape drives in robotic libraries, or upgrade

to another edition. See the Backup Exec licensing documentation for more information.

Refer to the Backup Exec hardware compatibility list

1. The hardware compatibility list (HCL) can help you verify that Backup Exec supports your devices, inquiry strings, and connection methods. Refer to the HCL to confirm that the device is officially supported by Backup Exec.
2. Refer to the HCL to confirm that the inquiry string for the device matches the string in the HCL.
3. Refer to the HCL to confirm that the connection method that is used to connect the device to the server is supported. Many devices support multiple connection methods. These devices are listed more than once in the HCL. Some connection methods may work for a particular device, but might not be listed in the HCL. If a method is not listed in the HCL, it is not officially supported.

Verify admin rights for the Backup Exec service account

1. Verify that the Backup Exec service account is a Domain administrator account or a built-in administrator account.

You can use the VQA tool for this procedure.

- For Backup Exec 2012 and later, click on the Backup Exec button, then on **Configuration and Settings**, then **Backup Exec Services**.
2. Verify that the account has the following basic rights and permissions:
 - Act as part of the operating system
 - Backup files and directories
 - Create a token object
 - Log on as a batch job
 - Log on as a service
 - Manage auditing and security log
 - Restore files and directories
 - Take ownership of files and other objects

Stop and disable the Removal Storage service

1. This procedure only applies to Windows Server 2003/XP. In Windows, click **Start**, then **Control Panel**, then **Administrative Tools**, and double-click on **Services**.

You can use the VQA tool for this procedure.

2. Click **Removal Storage Service**, then click on **Stop** and **Disable**. If a Removal Storage service error appears, disregard it. The error should not occur again after you restart the server.
3. Restart the server.

Review system event logs for hardware errors

1. In Windows, right-click on **My Computer**, then on **Manage**.
2. Expand **Computer Management > System Tools > Event Viewer**, then click on **System**. On some versions of Windows, the path is **Computer Management > System Tools > Event Viewer > Windows Logs > System**.
3. Look for any hardware errors in the logs.

Uninstall any third-party applications that control the tape device

1. Backup Exec may not be able to communicate properly with tape devices if another vendor's backup application is installed. This can happen even if the services for that application are disabled. Click **Start > Control Panel > Add/Remove Programs**.

Note: Do not uninstall tape diagnostic tools when Backup Exec services are running. Remove tape diagnostic tools after the diagnostic tests are complete.

2. Select the third-party application or tool, and click **Uninstall**.
3. Repeat the uninstallation process for all third-party applications or tools which prevent Backup Exec from communicating properly.

Use the Discover tool to troubleshoot hardware errors or reservation conflicts

1. The Discover tool displays detailed attributes of any backup devices that are attached to the server. To begin, stop the Backup Exec services.
2. Navigate to the Backup Exec installation directory.
3. To run the tool and create a text output, run the following command at a command prompt: `Discover.exe > C:\discover.txt`

Use tracer.exe to troubleshoot SCSI issues

1. Begin by verifying that the SCSI changer service is running. Click **Start**. In the search box, type `msinfo32` and press Enter.

You can use the VQA tool for this procedure.

2. In the **System Information** dialog, expand **Software Environment**, and then click **System Drivers**.
3. Locate the **SCSIChanger** device, and verify that its **Status** is set to **Running**.
4. After you have verified that the SCSI changer service is running, stop the Backup Exec services.
5. Navigate to your Backup Exec directory, and launch `tracer.exe`. The tracer program begins capturing SCSI information.
6. Restart the Backup Exec Services. To restart the services, launch `Servicesmgr.exe` from the Backup Exec directory.
7. After the services start, review the tracer log for any hardware errors or reservation conflicts.

Detect and uninstall any orphaned devices

1. Orphaned devices may be present in Device Manager after you replace hardware, or perform firmware and driver updates. To begin, open the Windows Device Manager using a command prompt by running the following commands at the command line:

```
C:\>cd WINDOWS
```

```
C:\WINDOWS>cd system32
```

```
C:\WINDOWS\system32>devmgmt.msc
```

```
C:\WINDOWS\system32>
```

2. In the Device Manager, click the **View** menu, then select **Show Hidden Devices**. This option shows all device drivers, including those that are not currently installed and running on the computer.
3. Expand the following devices. For any of these devices that are not shown in bold, right-click them and click Uninstall:
 - Medium changer
 - Tape drives
 - SCSI cards

Devices that are not bold are not loaded and can be uninstalled.

Warning: Removing devices and drivers that are still required by the system may result in the system becoming unstable and unable to start.

Disable, delete, and turn on the device in Backup Exec

1. In the Backup Exec console, go to Devices/Storage tab.
2. Right-click the device, and then select **Disable**.
3. Right-click the device, and then select **Delete**. A prompt to move the backup jobs to other devices or device pools may appear. For Backup Exec 2012, this prompt appears as **Retarget Jobs**.
4. Note the jobs that are affected and move the jobs back to the original device after resolving the issue with the device.
5. Stop the Backup Exec services.
6. Navigate to the Backup Exec directory and run `Tapeinst.exe`.
7. In the **Backup Exec Device Driver Installed** dialog, select the following options:
 - **Use tape drivers for all supported tape devices**
 - **Delete entries for tape devices that are unavailable, removed, or turned off**
 - **Use Plug-and-Play drivers for Windows 2000 and later**
8. Restart the Backup Exec services.

How to get more information about alerts and error messages

Backup Exec generates an error message when a condition occurs that is important enough to warrant your attention, or requires that you submit a response. Most alerts and error messages are self explanatory, but there may be times when you need to get more information to resolve a condition.

You can get more information on Backup Exec alert and error messages in the following ways:

- On the alert message, click the link for more information, or look in the job log and click the UMI link. This code is a hyperlink to the Arctera Technical Support website. You can access the technical notes that are related to the alert.
See [“Linking from the job log to the Arctera Technical Support website”](#) on page 282.
- Search the Arctera Technical Support knowledge base for the error.
See [“Searching the Arctera Knowledge Base”](#) on page 868.

Troubleshooting backup issues in Backup Exec

If you have problems with backing up data, review the following questions.

Table 20-2 Backup questions

Question	Answer
I am unable to back up certain files on my system that are in use by other processes. Why is that?	For non-snapshot backups, when Backup Exec encounters a file that is in use by another process, it either skips the file or waits for the file to become available. These actions depend on the options for no-snapshot backups that you configure when you create the backup. See “Configuring file and folder options for backup jobs” on page 683. If you configure Backup Exec to back up open files with a lock, it attempts to open the files in a different mode. It locks these files during backup to prevent other processes from writing to them. It is recommended that you close the applications that leave files open so that the files are backed up in a consistent state. To back up open files on Windows computers, use the Advanced Open File options to configure the backups that use snapshot technology. See “Configuring Advanced Open File options for backup jobs” on page 669.
Why does the Backup Exec Administration Console continue to own a storage device even when it's not running?	Backup Exec is a client/server application that must always be available to process the jobs that are submitted from both local and remote administrative consoles. The Backup Exec services claim all of the storage devices that are attached to the Backup Exec server whenever the services are running. Backup Exec requires constant control of the storage devices to collect statistics on media and storage device usage, and to provide media overwrite protection when necessary.

Table 20-2 Backup questions (*continued*)

Question	Answer
When I run a local backup, the total number of bytes backed up by Backup Exec does not match the number of bytes displayed by Windows. Why?	The type of partition for which the system is formatted may cause this problem. If you have a Windows NTFS compressed partition, Backup Exec displays the uncompressed byte count of the files that are backed up. Meanwhile, Windows Explorer displays the compressed byte count of the files on the hard drive. For example, Windows compresses an NTFS partition that contains 1 GB of data to 500 MB. Backup Exec reports that 1 GB of data was backed up, even though Windows Explorer displays that only 500 MB of compressed data exists on the hard drive If you have a FAT partition, Backup Exec reports the actual number of bytes of the files being backed up while File Manager reports an inflated amount of disk space. For example, a 2 GB FAT partition has a 32-K cluster size and File Manager displays 1.9 GB of used space. Backup Exec reports that 1.4 GB of data was backed up. Assuming that a 50-MB pagefile.sys is excluded from the backup, there is a 450-MB difference in the number of bytes. Converting to NTFS regains disk space since it is more efficient and the default cluster size (automatically set by Windows) in NTFS is less than FAT. Windows lets you specify a cluster size other than the default; however system performance may decrease. For more information, see the Windows documentation.
How can I change my existing database encryption key if I feel it has been compromised or to comply with company policies?	You can refresh a database encryption key at any time. Refer to the following topic for more information about refreshing database encryption keys. See “Refreshing Backup Exec Database encryption keys” on page 707.

Table 20-2 Backup questions (*continued*)

Question	Answer
Why do I experience slow throughput when I back up remote disks?	Local disk drives on the Backup Exec server can usually be backed up at a faster speed than backing up remote servers across a network. The backup speed for a remote disk is limited by the speed of the physical connection. The rate at which a remote server's hard disks are able to be backed up depends on the following items: ■ The make/model of network cards. ■ The mode/frame type configuration for the adapter. ■ The connectivity equipment (hubs, switches, routers, and so on). ■ The Windows settings.

Troubleshooting failed components in the SAN

Various problems can occur at any location in a SAN.

For Backup Exec to work properly, a storage device has to be recognized in the following locations:

- The bridge or router must recognize it as a SCSI device
- The operating system must recognize it as a device
- Backup Exec must recognize it as a supported device

In some cases, hardware issues may require you to contact your hardware vendor for technical support.

You may need to replace a component of your SAN, such as a bridge or a switch. For specific steps for replacing your equipment, refer to your hardware vendor's documentation.

See ["Troubleshooting offline storage devices in a SAN "](#) on page 860.

Troubleshooting offline storage devices in a SAN

If a device in your SAN has gone offline, follow these steps to determine the source of the problem.

Before you begin troubleshooting, verify that your storage devices are on the Backup Exec supported device list.

You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

Also, verify that all hardware drivers are up to date and are started. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table 20-3 Troubleshooting offline storage devices in a SAN

Step	Action
Step 1	Use the Windows Device Manager to verify that the operating system recognizes the device. If the device is not recognized, you may need to troubleshoot the device. See “Finding hardware errors in a SAN” on page 862.
Step 2	For robotic libraries, verify that robotic library support is installed.
Step 3	Check the system event log for the following errors, which indicate SAN communication errors: SCSI errors 9, 11, and 15, or timeout errors relating to storage. Check the application event log for multiple events 33152. These events indicate SAN communication errors. See “Finding hardware errors in a SAN” on page 862. You may need to contact your hardware vendor.
Step 4	If the library is online, but some or all of the drives are offline, use Backup Exec to initialize the library. See “Initializing a robotic library” on page 564.
Step 5	If initializing the library does not bring the storage devices online, check the library for an error display on the front panel, mechanical problems, or tapes inappropriately in the drives. Correct any errors that you find.

Table 20-3 Troubleshooting offline storage devices in a SAN (*continued*)

Step	Action
Step 6	If no errors exist on the library or if you corrected the errors and the storage devices are still offline, stop Backup Exec services and then restart. See “Starting and stopping Backup Exec services” on page 771.
Step 7	If restarting the services does not bring the storage devices online, restart the operating system. Be sure that no Backup Exec jobs are running when you restart.
Step 8	If restarting the operating system does not bring the storage devices online, reset the SAN to help identify problem tape storage. Recycling the SAN may also resolve Fibre Channel problems. See “Resetting the SAN” on page 863.

Finding hardware errors in a SAN

Use the following steps to find the common hardware errors that occur in a SAN. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table 20-4 Finding hardware errors in a SAN

Step	Action
Step 1	Verify that the proper device drivers are installed.
Step 2	Verify that the fibre cable is securely connected to the HBA and to the fibre switch.
Step 3	Verify that the fibre connection is properly connected to the robotic library from the fibre switch.

Table 20-4 Finding hardware errors in a SAN (*continued*)

Step	Action
Step 4	Check for a failed hardware component between the server and the fibre switch. Sometimes some of the servers in the SAN recognize the storage device, but other servers do not. If none of the servers in the SAN recognize the storage device, check for a failed hardware component between the fibre switch and the storage device.
Step 5	Reset the SAN, which may identify problem hardware components and may resolve fibre problems. See “Resetting the SAN” on page 863.

Resetting the SAN

Resetting the SAN involves turning off the components of the SAN, and then powering them on in a specific order.

Table 20-5 Resetting the SAN

Step	Action
Step 1	Turn off all servers, robotic libraries, and fibre bridges in the SAN. In rare cases, you may need to turn off the fibre switch also. If you need to turn off the switch, you should power it on before any other components. Wait for all checks to complete before you turn on the other components.
Step 2	Turn on the robotic library. See “Initializing a robotic library” on page 564.
Step 3	Verify that the fibre switch recognizes the robotic library.
Step 4	Turn on the central administration server.
Step 5	Verify that the operating system recognizes the robotic library and its drives.

Table 20-5 Resetting the SAN (continued)

Step	Action
Step 6	Turn on one of the managed Backup Exec servers. Wait for the managed Backup Exec server to start before you turn on the other managed Backup Exec servers.

Bringing storage devices online after an unsafe device removal event in a SAN

If a storage device is in use by Backup Exec at the time of an unsafe device removal event, the device appears as offline in Backup Exec.

Table 20-6 How to bring a device online after an unsafe device removal event

Step	Action
Step 1	Verify that no Backup Exec jobs are running in the SAN.
Step 2	Use Backup Exec to initialize the robotic library if the library is online but the drives are offline. See “Initializing a robotic library” on page 564.
Step 3	Stop all Backup Exec services and then restart them if the library is offline or if the drives are offline after initialization. If the device is not online, you may need to troubleshoot the device. See “Finding hardware errors in a SAN” on page 862.

Troubleshooting installation issues in Backup Exec

If you have problems with .NET installation review the following points.

- If the .NET 4.8 installation fails, try to install it manually before you retry installation of Backup Exec. This helps to isolate any issues from Backup Exec such as, issues with Microsoft's .NET installer or any environmental issues.
- The BKUPINST25.htm file logs any errors that occur during the .NET installation. However, detailed logs are available in .Net 4.8 setup logs, which

are available at %temp% folder and the file name are Microsoft .NET Framework 4.8 Setup_*.html, Microsoft .NET Framework 4.8 Setup_*.txt, and dd_ndp48-x86-x64-allos-enu_decompression_log.txt.

Troubleshooting blocked access to backed up items with GDPR Guard

If you have problems with blocking access to backed up items with GDPR Guard, review the following points.

- If the import command is successful but the restore and search view continue to display and restore the blocked items, check the following:
 - Check if the full path of any blocked item entry in the CSV file that was specified during import contains a comma (.). If the file name or any of the folders in the path contain a comma, the entry is not imported properly. For example: E:\Folder,Delimited\file.txt and E:\F1\ab,a.txt are unsupported entries for blocked items.
 - Some entries may get blocked if the full path of a blocked item's entry contains a wildcard within a folder that is also applicable to the items of that folder and its subfolders. For example: E:\F1*.txt blocks all text files starting with letter 'a' inside the E:\F1 folder and inside the E:\F1\F2 folder. To ensure that the item within a subfolder is not blocked, you must add each entry separately and not use wildcards in this scenario.
 - If the same server is added in Backup Exec using multiple names, such as by the NetBIOS name, FQDN, or IP address, the items are blocked only for the server name that is specified in the blocked item entry in the CSV. To resolve this issue, add the blocked item entries via each of the server names separately in the input CSV for the import operation.
- If some blocked items are skipped during the import operation with the reason displayed as invalid item path, check the following:
 - The full path for the skipped item does not start with the volume or share name and starts with server name. For example: E:\folder1*.txt and TestShare\F1\b.txt are acceptable entries in the full path.
 - Only the file name in the item path contains '*' related wildcards. If the folder path contains the wildcard, the entry is skipped. For example: G:\Test2*\CatTrans.xml and G:*\results.txt are invalid entries.
- If Import operations terminate with the following error: "An error occurred while processing a blocked item. The internal list of blocked items does not exist".

Check whether the input CSV file path is correct and does not point to an invalid drive or network location.

- If the export operation terminates with the following error: “An error occurred while opening the internal list of blocked items”.
 Ensure that the specified output folder path is valid and does not point to an invalid drive or network location..

- If some of the blocked items are displayed in the restore or search view that may be due to the following reasons:

- An import operation is running for the CSV file containing some blocked items of the same resource.
- The integrity check failed for the internal list of blocked items for that resource or some problem occurred while reading from the internal list of blocked files for the restore or search view.

If you run the restore job, the restore job log displays the exact reason of the failure in reading from the internal list of blocked files.

- If a soft or hard link path is specified in the CSV, only the link is blocked. To block the folder that the link is pointing to, specify the actual path of the folder in the CSV.
- If you want to remove a blocked file entry from an already imported CSV file, do the following in the order listed:

- Run the export command.
- Remove the entry from the generated CSV.
- Delete the internal file containing the list of blocked items for the resource whose entry is removed. Remove the internal file from the following location: "data\BLFileInfo" folder under the Backup Exec installation directory.
- Run the import command again with the `ResetIfCorrupted` parameter.

```
Import-BEItemsToBlock - RESETIfCorrupted - CsvFilePath <CVS
file path> - ColumnNameForServerName <Name of 1st column> -
ColumnNameForBlockedItemPath <Name of 2nd column>
```

- If the restore and search view does not display results for a long time it can be because a large number of blocked items are added for that resource. In such a scenario, restore job can still be run at the folder or volume level to view the items that are restored and viewed.
- If a standalone media server that has blocked items needs to be converted into an MBES server, perform the following steps in the order listed:
 - Export the blocked items file.

- Delete the bin files from the `Data\BLFileInfo` folder under the Backup Exec installation folder.
- Convert the standalone server to an MBES server.
- Import the blocked items on CAS using the exported file that you created.

See [“About GDPR Guard”](#) on page 841.

How to improve Backup Exec's performance

To get the best performance from Backup Exec, you should review several factors:

- Data transfer path
- Backup Exec agent performance
- Network performance
- Backup Exec server performance
- Storage device performance

For more information on how to measure and tune the performance of these items, see the following sections:

See [“Improving backup performance in Backup Exec”](#) on page 142.

See [“Troubleshooting backup performance”](#) on page 145.

Accessing Arctera Online

You can access Arctera community forums, learn about training courses, and view Arctera websites

Table 20-7 Arctera Online menu items

Item	Description
Share Your Ideas	Connects you to the forum where you can post your ideas for improving Backup Exec.
Education Services	Provides the links to all Arctera Education training and custom learning services.
Backup Exec Tech Center	Provides the links to Backup Exec self-paced training modules.
Backup Exec Page	Provides the links to resources for Backup Exec.

Table 20-7 Arctera Online menu items (continued)

Item	Description
Arctera Home Page	Connects you to the Arctera website.

To access Arctera Online

- ◆ Click the Backup Exec button, select **Arctera Online**, and then select the appropriate menu item.

See “[Searching the Arctera Knowledge Base](#)” on page 868.

Searching the Arctera Knowledge Base

The Arctera Knowledge Base is a centralized location where you can find more information about Arctera products. The knowledge base contains information about how to install, upgrade, configure, and use Arctera products. It also contains information about requirements, best practices, and how to troubleshoot problems. The Arctera Knowledge Base is accessible from within Backup Exec.

Note: You must have an active Internet connection to access the Arctera Knowledge Base.

The knowledge base uses a keyword-based search technology. It focuses on the important keywords in a search and compares them to other search phrases to provide the best possible results. You can use Boolean search features and expression queries to provide search parameters. For best results, focus on a few keywords that best represent your question.

When you search the knowledge base, a new browser window launches and displays the search results.

To search the Arctera Knowledge Base

- 1 Do either of the following:
 - Click the Backup Exec button, select **Technical Support**, and then select **Search Knowledge Base**.
 - On the **Home** tab, in the **Support** group, ensure that **Technical Support** is checked. Then, in the **Technical Support** panel, click **Arctera Technical Support**.
- 2 Enter a keyword or phrase, and then click the search icon.

Contacting Backup Exec Technical Support

If you have tried to solve a problem, but still need a resolution, you can contact Arctera Technical Support for Backup Exec over the Internet or by phone.

To expedite the Technical Support process, do the following:

- Know your Backup Exec version and revision number.
- Use one of the diagnostic utilities that are included with Backup Exec to collect the information that technical support can use to diagnose your issue.

To contact Backup Exec Technical Support

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Backup Exec Technical Support**.

See [“Displaying the version information for Backup Exec”](#) on page 116.

See [“About Backup Exec diagnostic tools”](#) on page 869.

See [“Accessing Arctera Online”](#) on page 867.

Using Remote Assistance

Remote Assistance launches a WebEx session on the Internet, which lets you join a support session or start a support session.

To use Remote Assistance

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Arctera Remote Assistance**.

Managing your Backup Exec support cases

From Backup Exec, you can launch the Support website, where you can create, review, and manage cases related to technical product support.

To manage your Backup Exec support cases

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Manage Support Cases**.

About Backup Exec diagnostic tools

The following diagnostic tools help to troubleshoot issues in Backup Exec:

Table 20-8 Backup Exec diagnostic tools

Item	Description
Veritas QuickAssist (VQA) Help Tool	Scans the local computer and generates a report about common issues in your Backup Exec environment. See “Running the Veritas QuickAssist Help Tool” on page 870.
Backup Exec diagnostic application	Gathers the pertinent information about a Windows computer for troubleshooting. See “Generating a diagnostic file for troubleshooting Backup Exec” on page 871.
Gather utility for Linux servers	Creates and compiles a Packet file. The file contains detailed information about installation, diagnostics, and error reports. See “Running the begather utility to troubleshoot Backup Exec components on Linux servers” on page 875.
Backup Exec Debug Monitor	Captures the debug output from Backup Exec and saves it into debug logs. See “Using the Backup Exec Debug Monitor for troubleshooting” on page 875.

Running the Veritas QuickAssist Help Tool

The Veritas QuickAssist help tool is a multi-product diagnostic utility that identifies common issues, gathers data for support-assisted troubleshooting, and provides links to other customer self-help and support resources. To perform checks on the local computer, Veritas QuickAssist gathers information from the local computer and then analyzes it. Veritas QuickAssist does not permanently alter any files on the computer unless you select the option when you run the tool. Veritas QuickAssist does not permanently install anything on your computer when it runs.

To run the Veritas QuickAssist Help Tool

- 1 Click the Backup Exec button, select **Technical Support**, and then select **Quick Assist**.
- 2 Follow the on-screen prompts.

Generating a diagnostic file for troubleshooting Backup Exec

Backup Exec includes a diagnostic application (Bediag.exe) that gathers information about a Windows computer for troubleshooting purposes. This application can be run from within Backup Exec or it can be run from a command line. This application can be run for a local server or a remote server. You can run diagnostics on a remote Backup Exec server if Backup Exec is installed on the remote server and the Backup Exec services are running.

The Bediag application collects the following types of information:

- Account groups, account privileges, and environment settings.
- Backup Exec software version and registry information, Backup Exec agent listing, Windows version information, SCSI hardware configuration, SQL Server information, Driver services information, and Windows Services information.
- Server information, supported shared directories, and Windows sockets information.

To generate a diagnostic file from within Backup Exec

- 1 Click the Backup Exec button, select **Technical Support**, and then select **Backup Exec Diagnostics**.
- 2 Select the server from the drop-down list.
 If the server that you want to select is not in the drop-down list, click **Browse** to select it from a list of available servers.
- 3 Enter the logon information for the server that you selected.
- 4 Click **Run Diagnostics**.
- 5 Click **Close**.
- 6 Open the "Bediag.txt" from the directory that contains Bediag.exe (by default *<Backup Exec install path>\Backup Exec*).

To generate a diagnostic file from a command line

- 1 Launch the command prompt.
- 2 Do one of the following:

To generate a diagnostic file for a Backup Exec server

From the directory `<Backup Exec install path>\Backup Exec\`, type `bediag [switches] servername` .

See [“Command line switches for a diagnostic file”](#) on page 872.

To generate a diagnostic file for a remote computer

From the directory `<Backup Exec install path>\Backup Exec\`, type `bediag [switches] workstationname`.

See [“Command line switches for a diagnostic file”](#) on page 872.

- 3 Open the "Bediag.txt"from the directory that contains Bediag.exe (by default `<Backup Exec install path>\Backup Exec`).

Command line switches for a diagnostic file

You can add the following switches to gather additional information when you generate a diagnostic file for troubleshooting.

See [“Generating a diagnostic file for troubleshooting Backup Exec”](#) on page 871.

Table 20-9 Command line switches for a diagnostic file

Switch	Description
/a	Dumps the Agent List.
/ad	Dumps Active Directory information.
/adamm	Appends ADAMM file information.
/agents	Dumps agent information for the favorite resources.
/all	Dumps everything.
/app	Dumps the Application Event log.
/b2d	Dumps Backup Exec backup-to-disk information.
/basicscsi	Dumps basic SCSI hardware subkey from the registry.

Table 20-9 Command line switches for a diagnostic file (*continued*)

Switch	Description
/beallfiles	Dumps all Backup Exec files and directories.
/bedb	Dumps Backup Exec Database information.
/befiles	Dumps Backup Exec file information.
/bereg	Dumps Backup Exec software configuration from the registry.
/beupdate	Dumps Backup Exec update information.
/bex	Dumps only the Backup Exec entries that are in the Application Event log.
/c	Dumps the Backup Exec software configuration from the registry.
/caso	Dumps information about the central administration server and managed Backup Exec servers.
/cps	Dumps CPS registry information.
/detailnic	Dumps the server's detailed network card information.
/detailscsi	Dumps the server's detailed SCSI adapter information.
/dirsvc	Dumps the directory service event log.
/dlo	Dumps DLO ini file.
/dns	Dumps the DNS event log.
/err	Dumps only the error events from any event log.
/evdb	Dumps Arctera Enterprise Vault database information.
/evevents	Dumps the Arctera Enterprise Vault event logs.
/evreg	Dumps Arctera Enterprise Vault registry information.
/frs	Dumps the file replication service event log.
/o:[file]	Specifies the output job log for append. Omitting [file] sends output to the screen.
/h	Dumps the SCSI hardware subkey from registry.
/instapp	Dumps information about installed applications.
/lic	Dumps Backup Exec server license information.

Table 20-9 Command line switches for a diagnostic file (*continued*)

Switch	Description
/liveupdate	Dumps Arctera Update Information
/n	Dump Windows Socket Network Protocols.
/networkinformation	Dumps TCP/IP settings, Winsock information, and Windows firewall information.
/oracle	Dumps Oracle information.
/p	Dumps the user privileges.
/power	Dumps the PowerShell event log.
/recs:n	Dumps only the newest records from given event logs. The bex, err, and recs switches must be used with the app switch and the sys switches.
/s	Dumps the information on Services.
/sec	Dumps the security event log.
/server	Dumps server information, such as CPU, memory, disk information, and more.
/services	Dumps information about services.
/sql	Dumps the Microsoft SQL Server information.
/bereginfo	Dumps the Backup Exec registry information.
/sys	Dumps the system event log.
/sys32info	Dumps driver information from the system32\drivers folder.
/svcacct	Dumps user names used for Backup Exec services.
/u	Dumps Microsoft update information.
/userinfo	Dumps user privileges and group information.
/v	Dumps Server Information.
/w	Dumps Windows version information.
/winpower	Dumps the Windows PowerShell event log.
/winupdate	Dumps Microsoft update information.

Table 20-9 Command line switches for a diagnostic file (continued)

Switch	Description
/x	Dumps Microsoft Exchange Server Information.
/?	Displays usage information.

Running the begather utility to troubleshoot Backup Exec components on Linux servers

The begather utility brings together the files that help you diagnose issues with Backup Exec components on Linux servers. After you run it, the begather utility displays the name of the Packet file that it creates. The files that are gathered contain detailed information regarding installation, diagnostics, and error reporting. Reviewing these files before contacting technical support can reveal the source of the issue. If the solution is not evident based on the gathered files, have the Packet file available when contacting support. The support technician may request an email that contains the Packet file.

Run the begather utility to troubleshoot Backup Exec components on Linux servers

- 1 Log on as root to the Linux server on which the Backup Exec components are installed.
- 2 Navigate to the following directory:
`/opt/VRTSralus/bin`
For example:
`cd /opt/VRTSralus/bin`
- 3 Start the begather utility.
For example:
`./begather`
- 4 Note the location of the Packet file that is displayed on the screen.

Using the Backup Exec Debug Monitor for troubleshooting

The Backup Exec Debug Monitor, or SGMon, is a diagnostic tool that captures debug output from Backup Exec and saves it in debug logs. SGMon debug logs

can help you troubleshoot backup issues. Furthermore, debug logs can help Technical Support diagnose and repair problems.

When you open SGMon, it automatically captures debug data from Backup Exec's services. To collect debug information while SGMon is closed, enable debug log creation outside of SGMon and specify a directory in which to save the logs.

For more information about how to configure the Debug Monitor and read log files, refer to the help within the Debug Monitor.

To use the Backup Exec Debug Monitor for troubleshooting

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Collect debug output**.

About the Backup Exec debug tool

Backup Exec includes a debug tool (BEDBG) that generates diagnostic information about the Backup Exec processes that shut down unexpectedly. The diagnostic information helps Technical Support diagnose and repair the problem. The Backup Exec debug tool runs by default in Backup Exec. The data that the tool gathers is copied into the BEDBG folder which is located in *<Backup Exec install path>\Backup Exec*.

Simplified Disaster Recovery

This chapter includes the following topics:

- [About Simplified Disaster Recovery](#)
- [Requirements for using Simplified Disaster Recovery](#)
- [Preparing computers for use with Simplified Disaster Recovery](#)
- [How to ensure that backups are enabled for Simplified Disaster Recovery](#)
- [How Simplified Disaster Recovery uses disaster recovery information files](#)
- [Setting or changing the alternate location for the disaster recovery information file](#)
- [Creating a Simplified Disaster Recovery disk image](#)
- [Preparing to recover from a disaster by using Simplified Disaster Recovery](#)
- [Recovering a computer with Simplified Disaster Recovery](#)
- [Advanced Disk Configuration on the Recover This Computer Wizard](#)
- [Performing manual disaster recovery](#)

About Simplified Disaster Recovery

Simplified Disaster Recovery (SDR) is automatically installed with Backup Exec so that you can perform disaster recovery on Windows computers on which the Agent for Windows is installed. By default, Backup Exec is configured to back up all of the critical system components that you need to perform a disaster recovery.

After a computer's critical system components are backed up, use the **Create Simplified Disaster Recovery Disk Wizard** to create a Simplified Disaster Recovery disk image. Use the recovery disk to perform disaster recovery of the computers that are backed up.

When you use SDR to perform a recovery, Backup Exec uses the system-level information from the SDR backup to rebuild the server and restore it to a functional state. Recovery includes a bare metal or dissimilar hardware restore operation.

Through integration with the Microsoft Volume Shadow Copy Service (VSS), the SDR backups include all selected elements, even if they are components of the active operating system or are in an open state. VSS integration ensures that the backups are in a consistent state and have been properly placed into a quiescent state when the backup occurs.

Simplified Disaster Recovery is available only for servers on which the Agent for Windows is installed and that are backed up through the Agent for Windows. You must purchase the Agent for Windows separately, and then install it on the remote computers that you want to protect.

The Simplified Disaster Recovery feature now requires the Microsoft Assessment and Deployment Kit (ADK) to create a Simplified Disaster Recovery disk image (.iso).

Note: After the Backup Exec upgrade, customers must customize the existing SDR ISOs to make the ISO compatible with the Backup Exec release to which they have upgraded.

For information about the best practices to use Backup Exec Simplified Disaster Recovery (SDR), refer to *Backup Exec Best Practices*.

See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 881.

See [“How to ensure that backups are enabled for Simplified Disaster Recovery”](#) on page 885.

See [“Requirements for using Simplified Disaster Recovery”](#) on page 878.

See [“Creating a Simplified Disaster Recovery disk image”](#) on page 892.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 916.

Requirements for using Simplified Disaster Recovery

The following items are required for Simplified Disaster Recovery (SDR):

- Backup Exec or the Backup Exec Agent for Windows must be installed on any computers that you want to protect with SDR.

Note: Backup Exec only supports 64-bit Simplified Disaster Recovery disk image. This recovery disk image can be used to recover both 32-bit and 64-bit systems.

- Encryption key files must exist for all volumes that you encrypt with Windows BitLocker Drive Encryption.
- A third-party ISO 9660-compliant CD or DVD burning application must be available to burn the SDR-created bootable image to a CD or DVD.
- A writable or rewritable CD or DVD device must be available.
- An Internet connection so that you can download the Microsoft Assessment and Deployment Kit.
- The option **Use storage-based catalogs** must be enabled. If you disable this option, the backup sets that you create for use with SDR cannot be restored during an SDR recovery operation. As a result, SDR cannot recover the failed computer. To ensure that this option is enabled, click the Backup Exec button, click **Configuration and Settings > Backup Settings > Catalog**.

Note: If you use deduplication disk storage devices, be aware that there are limitations in their use with SDR.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature”](#) on page 924.

Additional requirements exist when you create a Simplified Disaster Recovery disk image and when you run the **Recover This Computer Wizard**, as described in the following sections.

Requirements for creating a Simplified Disaster Recovery disk image

The following items are required to create the Simplified Disaster Recovery disk image:

- The Simplified Disaster Recovery disk image must be the same version of Backup Exec as the Backup Exec server. You cannot use SDR to restore the backups that were created with previous versions of Backup Exec.
- From Windows Server 2012 to Windows Server 2016, you can create an SDR disk only using ADK 10. For Windows Server 2019, you must use

Windows ADK 1809 to create an SDR disk. If the Backup Exec server does not run on Windows Server 2012 or later, or if the server does not have an Internet connection, methods are provided to let you create the Simplified Recovery Disk.

Note: The recovery disk created with Windows ADK 10 does not detect the Backup Exec Storage folder (\BEData) if this folder is configured on the storage pools and storage spaces on the Backup Exec Server. This problem happens only when the Backup Exec server installed on Windows Server 2019 operating system is not available and you want to perform an SDR Local recovery from the BEData folder.

If you create the storage pools and spaces using the recovery disk created with Windows ADK 10, then after system restore, the Windows Server 2012 and Windows Server 2012 R2 operating systems does not detect the storage pools and spaces.

If you create the storage pools and spaces using the recovery disk created with Windows ADK 1809, then after system restore, all Windows 2016 and earlier operating systems do not detect the storage pools and spaces.

- 5 GB of disk space to download and install ADK.
- 1 GB of disk space to store the required files and folders to create the ISO image.

Requirements for running the Recover This Computer Wizard

The following items are required to run the **Recover This Computer Wizard**:

- The Simplified Disaster Recovery disk image must be the same version of Backup Exec as the Backup Exec server.
- If the computer that you want to recover was backed up to a tape device, deduplication storage, or to a virtual disk, then SDR cannot store the disaster recovery information file with the backup sets. Instead, you must provide the path to the default location or to the alternate location when the **Recover This Computer Wizard** prompts you. The default location is `C:\<Backup Exec install path>\Backup Exec\sdr\Data`. If the file is unavailable, you cannot recover the computer with SDR.
- The backup set that contains all of the critical system components for the computer that you want to restore.
- The boot drive on the computer that you want to recover must have from 3-GB to 5-GB of free space depending on the operating system and configuration.

If a blank screen appears and the computer does not restart after you use the SDR disk, ensure that the boot drive has the necessary amount of free space. Then, restart the computer again.

- The target computer that you want to recover to must have an amount of RAM that is equal to or greater than the minimum that the restored operating system requires.
- The target computer that you want to recover to must have disks with enough free space to contain the data that you restore from the backup of critical volumes. A volume is considered critical if it is required for the computer to start successfully.
- If you recover a Windows computer that has BitLocker encryption enabled, you must turn on BitLocker encryption after the restore. See your Microsoft documentation for more information on BitLocker drive encryption.
- If the computer you recover contains a RAID setup, you may be required to configure the RAID before you start it with the SDR disk. Use the computer manufacturer's RAID software to configure the RAID system.
- If you restore Windows with storage spaces and storage pools, you should be aware of the possible restore scenarios.
See [“Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces”](#) on page 922.

See [“Creating a Simplified Disaster Recovery disk image”](#) on page 892.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 916.

Preparing computers for use with Simplified Disaster Recovery

Prepare computers for use with Simplified Disaster Recovery (SDR) by performing the following steps:

Table 21-1 Preparing computers for use with Simplified Disaster Recovery

Step	Description	More information
Step 1	Specify an alternate location where copies of the disaster recovery information files are stored.	Backup Exec creates the disaster recovery information file after a backup job that includes all critical system components completes successfully. Backup Exec then stores the disaster recovery information file in the default storage location with the backup set on a disk storage or disk cartridge device, and in the alternate storage locations. Catalog entries from subsequent backups are automatically added to the disaster recovery information file. Warning: If you back up a computer to a tape device, deduplication storage, or a virtual disk, Backup Exec cannot store a disaster recovery information file with the backup set. You must have a disaster recovery information file in an alternate location to ensure that you can use SDR to recover the computer. See “Setting or changing the alternate location for the disaster recovery information file” on page 889.

Table 21-1 Preparing computers for use with Simplified Disaster Recovery
(continued)

Step	Description	More information
Step 2	Run the backup jobs that include all critical system components (SDR-enabled backups) for the computers that you want to protect. These are the backups for which the Simplified Disaster Recovery indicator is ON. Note: For environments running the Central Admin Server feature, run a database maintenance job before you run SDR-enabled backups. Otherwise, the central administration server denies communication attempts from the managed Backup Exec servers. See “Configuring database maintenance and security” on page 701.	By default, Backup Exec selects all critical system components when you select a computer for backup. When all critical system components are included in the backup job selections, the Simplified Disaster Recovery indicator on the backup selections appears as ON. If you deselect one or more critical system component files, the indicator appears as OFF. It is recommended that you select the entire computer for backup; otherwise, Backup Exec cannot create the system-specific disaster recovery information file. Critical system components include the following: ■ System volume (including EFI and utility partitions) ■ Boot volume (executing operating system) ■ Services application volumes (boot, system, and automatic startup) ■ System State devices and volumes (including Active Directory, system files, and so on) See “How to ensure that backups are enabled for Simplified Disaster Recovery” on page 885.

Table 21-1 Preparing computers for use with Simplified Disaster Recovery
(continued)

Step	Description	More information
Step 3	Create additional copies of the disaster recovery information files and store them in a safe place.	Backup Exec stores the important disaster recovery information files in the default path, the alternate location, and with the backup set if disk storage or disk cartridge devices are used as the destination storage. It is recommended that you also make additional copies of the files and store them in a safe place. Without the disaster recovery information files, you cannot recover Backup Exec servers by using Simplified Disaster Recovery (SDR). Having multiple copies of the disaster recovery information files ensures that you can successfully recover Backup Exec servers with SDR. By default, the disaster recovery information files reside in the <i><Backup Exec install path>\Backup Exec\SDR\Data</i> directory on the Backup Exec server. Use Windows Explorer or another Copy utility to copy the disaster recovery information files from the default location to another storage location of your choice. Backup Exec does not automatically update these copies, but they do let you restore a computer to an older point-in-time if the other disaster recovery files are not available.
Step 4	Use the Create Simplified Disaster Recovery Disk Wizard to create a disk image, and then a bootable CD or DVD recovery disk.	Backup Exec generates alerts after each backup job until you create the Simplified Disaster Recovery disk image. You can disable these alerts, but it is recommended that you create the Simplified Disaster Recovery disk image. See “Creating a Simplified Disaster Recovery disk image” on page 892. If you install the Create Simplified Disaster Recovery Wizard on a standalone server, you cannot use the wizard to run Arctera Update.

How to ensure that backups are enabled for Simplified Disaster Recovery

By default, Backup Exec is configured to back up all of the critical system components that you need to use Simplified Disaster Recovery to recover a computer.

When the Simplified Disaster Recovery indicator for the backup selections is green, or **ON**, it indicates that the critical system components are selected for backup. The backup is SDR-enabled. If the indicator is gray, or **OFF**, then the backup is not enabled for SDR. You can click the icon to select all of the necessary components for disaster recovery, or to disable disaster recovery for backups created by this job.

To ensure that a backup is enabled for Simplified Disaster Recovery, you can view the backup selections that appear when you create or edit a backup job.

To ensure that backups are enabled for Simplified Disaster Recovery

- 1 On the **Backup and Restore** tab, create a new backup job or edit an existing one.
- 2 On the **Backup Definition Properties** dialog box, on the selections pane, do one of the following:
 - Ensure that the icon to the left of the computer name is green.
 - Click **Edit**, and ensure that the icon on the right is green, and that the text indicates that Simplified Disaster Recovery is on.

Note: If you hover the mouse over the icon, the text indicates whether the Simplified Disaster Recovery is on or off.

- 3 If the icon is not green, or if the text indicates that Simplified Disaster Recovery is off, click the icon and then select the option **Select all necessary components for disaster recovery**.

Figure 21-1 Simplified Disaster Recovery indicator is ON in the Backup Definition Properties dialog box





Figure 21-2 Simplified Disaster Recovery indicator is ON in the Backup Selections dialog box



See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 881.

How Simplified Disaster Recovery uses disaster recovery information files

For each computer that you back up and for which the Simplified Disaster Recovery indicator is **ON**, Backup Exec creates a disaster recovery information file. A disaster recovery information file contains computer-specific information for the computer that is backed up. Each time that a backup of all critical system components runs, the disaster recovery information files are automatically updated. Each disaster information recovery file uses the file name <computer_name>.DR. SDR uses the computer-specific information that is contained in the file when you run the **Recover This Computer** wizard. Without a disaster recovery information file, a recovery of the computer is not possible with SDR.

Note: Backup Exec by default supports the latest three full SDR backup chains required for system recovery using SDR. Each backup chain includes one full backup set, its dependent incremental and differential backups, and their duplicate backup sets.

A disaster recovery information file contains the following information for the computer that is backed up:

- Hardware-specific information, such as hard disk partition information, mass storage controller information, and network interface card information.
- A list of catalog entries that identify the backup sets and storage media that are needed to recover the computer.
- The Windows Automated System Recovery configuration information file (asr.xml) required during the recovery process.

Backup Exec stores the disaster recovery information file in the following locations:

- With the backup sets if the backup storage is disk storage or a disk cartridge device.
- On the Backup Exec server's hard drive in the following path:

```
C:<Backup Exec install path>\Backup Exec\sdr\Data\
```

- In an alternate location that you specify, on a different computer than the Backup Exec server.

Note: It is recommended that you specify an alternate storage location. If the Backup Exec server crashes, you cannot retrieve the disaster recovery information file from the default location. However, you can retrieve it from the alternate location. You should also create additional copies of the disaster recovery information files and store them in a safe place. Use Windows Explorer or another Copy utility to copy the disaster recovery information files from the default location to another storage location of your choice. Backup Exec does not automatically update these copies, but they do let you restore a computer to an older point-in-time if the other disaster recovery files are not available.

If the disaster recovery information file is stored with the backup sets, then SDR automatically uses that file to perform the recovery. If the backup sets are stored on a tape storage device, deduplication storage, or on a virtual disk, then SDR cannot store the file with the backup sets. Instead, you must provide the path to the default location or the alternate location of the disaster recovery file when the **Recover This Computer** wizard prompts you.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 889.

See [“Changing the default path for the disaster recovery information files”](#) on page 890.

Setting or changing the alternate location for the disaster recovery information file

It is recommended that you specify an alternate location where Backup Exec can store the disaster recovery information files. The disaster recovery information files contain computer-specific information for each computer that you back up with Simplified Disaster Recovery (SDR). When you recover a Backup Exec server from a disaster, you must have the Backup Exec server's disaster recovery information file available. Without it, you cannot use SDR to recover the Backup Exec server.

It is recommended that the alternate location be on another computer or on a different physical drive than the default location. If the Backup Exec server's hard drive is damaged, you can access a copy of the disaster recovery information from the alternate location. You can specify a drive letter that is mapped to a network share as the alternate location.

To use a remote computer's hard drive as the alternate path, establish a valid connection to the remote computer. Specify a UNC path as the alternate path, and then check the directory to ensure that the disaster recovery information files are copied.

For a remote Backup Exec server, specify an alternate location on a shared drive. You should also create additional copies of the disaster recovery information files and store them in a safe place.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 887.

To set or change the alternate location for the disaster recovery information file

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Simplified Disaster Recovery**.
- 3 In the **Alternate path** field, enter the path where you want to store a copy of the disaster recovery information file, or click **Browse** to navigate to a path.

4 Specify the logon account to use to access the alternate path.

5 Click **OK**.

See [“Changing the default path for the disaster recovery information files”](#) on page 890.

Changing the default path for the disaster recovery information files

You can change the default path for the disaster recovery information files. However, it is recommended that you do not change the default path.

Copies of the disaster recovery information files are necessary to automate the recovery of a Backup Exec server.

Backup Exec automatically creates the disaster recovery information file during a backup and stores a copy of it in the following path:

```
C:<Backup Exec install path>\Backup Exec\sdr\Data\
```

If you change the disaster recovery data path, ensure that you copy the existing disaster recovery information files to the new path. The copies let you recover a computer with the previous SDR backups.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 887.

To change the default path for the disaster recovery information files

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Simplified Disaster Recovery**.
- 3 In the **Path** field, change the path to the location where you want to store the disaster recovery information file, or click **Browse** to navigate to a location.
- 4 Click **OK**.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 889.

Disaster recovery information file data paths

Simplified Disaster Recovery (SDR) creates the disaster recovery information files that are required for recovery.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 887.

Table 21-2 Storage locations for the disaster recovery information file

Item	Description
Path	Indicates a directory path where you want to store the disaster recovery information files for the computers that you back up. Backup Exec automatically creates the files after a backup job runs for which all critical system components are selected. Backup Exec then stores the disaster information recovery files in the following default location: <i>C:<Backup Exec install path>\Backup Exec\sdr\Data\<computer name>.dr.</i> Note: It is recommended that you do not change the default path.
Alternate path	Indicates an alternate path where you can store copies of the disaster recovery information files for the computers that you back up. Backup Exec automatically creates or updates the files after a backup job runs for which all critical system components are selected. Backup Exec then stores copies of the disaster recovery information files in this location. It is recommended that you specify an alternate path that is not on the Backup Exec server, or that is on a different physical drive than the default location. You can specify a drive letter mapped to a network share as the alternate location. If the Backup Exec server's hard drive is unavailable during a recovery, you can copy the disaster recovery information file from the alternate path to any location to recover the failed computer. To use a remote computer's hard drive as the alternate path, establish a valid connection to the remote computer. Specify a UNC path as the alternate path, and then check the directory to ensure that the disaster recovery information files are copied.

Creating a Simplified Disaster Recovery disk image

The **Create Simplified Disaster Recovery Disk Wizard** guides you through the process of creating a startup recovery disk image that you can use to recover backed up computers.

The following items are required to create the Simplified Disaster Recovery disk image:

- 5 GB of disk space to download and install the ADK.
- 1 GB of disk space to store the required files and folders to create the ISO image.

Warning: The ADK download may take few hours depending on your network speed.

After you run Simplified Disaster Recovery-enabled backups of the computers that you want to protect, run the **Create Simplified Disaster Recovery Disk Wizard** to create the Simplified Disaster Recovery disk image. Simplified Disaster Recovery-enabled backups are those backups that display the green indicator on the backup properties pane that indicates the critical system components on the computer are selected.

See [“How to ensure that backups are enabled for Simplified Disaster Recovery”](#) on page 885.

For Windows Server 2012 to Windows Server 2016, the wizard installs the ADK. For Windows Server 2019 and Windows Server 2022, you must manually download and install the ADK.

The wizard uses the required files and folders from the ADK to create a startup recovery disk image in ISO format. You can then burn the image to a CD or DVD by using a third-party image burning application. For disaster recovery of a computer, you use the Simplified Disaster Recovery disk to start the computer, and then to recover it.

When you create the recovery disk, you can add language and time zone selections. You can also select the computers for which you want to add network and storage drivers. You can further customize the recovery disk by adding new OEM drivers.

Note: If you create the Simplified Disaster Recovery disk image before you run backups, the drivers for the backed-up computers are not included on the recovery disk. You must create a new recovery disk if you want the drivers of the backed-up computers to be included on the recovery disk. You can then specify the existing recovery disk as the source, and then select the computers for which you want to add the drivers to the recovery disk.

Depending on your environment, different methods are available for you to create the Simplified Disaster Recovery disk image.

Table 21-3 Methods for creating the Simplified Disaster Recovery disk image

Environment	Method
If the Backup Exec server runs on Windows Server 2019 and Windows Server 2022	See the section called “Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2019 and Windows Server 2022” on page 893.
If the Backup Exec server runs on Windows Server 2012 to Windows Server 2016	See the section called “Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2012 to Windows Server 2016” on page 897.
If the Backup Exec server is not available	See the section called “Creating a Simplified Disaster Recovery disk image if the Backup Exec server is not available” on page 900.
If the Backup Exec server on which you want to create the Simplified Disaster Recovery disk image does not have an Internet connection and does not have ADK installed	See the section called “Creating a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an internet connection and does not have ADK or Windows Preinstallation Environment installed” on page 903.

Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2019 and Windows Server 2022

If the Microsoft Windows Assessment and Deployment Kit (ADK) is not already installed on this server, you have to manually download and install Windows ADK 1809 or Windows Server 2022 compatible ADK using the following link:
<https://docs.microsoft.com/en-us/windows-hardware/get-started/adk-install?ocid=ti-235208000>

It is recommended that you download and install the required Windows ADK version before you launch the Create Simplified Disaster Recovery Disk Wizard.

If Windows ADK version is already installed, the Create Simplified Disaster Recovery Disk Wizard detects the ADK installed and proceeds to create the recovery disk.

To create a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2019 and Windows Server 2022

- 1 Do any of the following:
 - On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 2 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 3 Do one of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

- 4 Click **Next**.

5 Do one of the following:

If the Windows ADK 1809 is already installed on the computer The Simplified Disaster Recovery Disk Creation wizard detects the installed version of ADK and the wizard continues with the recovery disk creation process.

Note: If you have created an SDR using ADK 10.0, you can customize the SDR even if you have Windows ADK 1809 installed on a Windows Server 2019 operating system.

If the Windows ADK 1809 is not installed on the computer The Simplified Disaster Recovery Disk Creation wizard displays the following link to manually download and install Windows ADK 1809.

<https://docs.microsoft.com/en-us/windows/hardware/get-started/adk-install#id=235208000>

On the Windows ADK 1809 page, install the **Windows ADK Insider** and **Windows Preinstallation Environment** components.

For the **Windows ADK Insider** component, on the Windows Assessment and Deployment Kit wizard, follow the prompts to install the ADK. On the **Select the features you want to install** page, select only the **Deployment Tools** check box.

After the ADK is installed, install the **Windows Preinstallation Environment** component by accepting the recommended defaults in the wizard.

After both components are installed, launch the Simplified Disaster Recovery Disk Creation wizard again.

If the ADK 10.0 is installed on the computer The Simplified Disaster Recovery Disk Creation wizard detects that ADK 10.0 version is installed and the ADK update screen is displayed.

The wizard informs you that this is not the latest available version of ADK but you can proceed with the SDR disk creation. Confirm that you want to continue using ADK 10. If you continue, it may not support all the features of Windows Server 2019 and Windows Server 2022.

After you select the check box and click **Next**, the wizard continues with the recovery disk creation process.

6 Specify the startup options that you want when you recover a computer, and then click **Next**.

7 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups

Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**.

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include in the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the **Select computers to use the drivers from** panel, click **Next**.

- 8 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: If you use any non-English characters in the volume label, then those characters are not properly displayed when you mount the .iso image.

Note: It is recommended that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows: `C:\<Backup Exec install path>\Backup Exec\sdr\CustomSDRImage.iso`

- 9 Review the summary panel, and then click **Create Image**.
- 10 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 11 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2012 to Windows Server 2016

If the Microsoft Windows Assessment and Deployment Kit (ADK) is not already installed on this server, then the Create Simplified Disaster Recovery Disk Wizard guides you through the download and installation of the ADK. The ADK installation defaults are already selected. It is recommended that you use these defaults.

Note: Backup Exec supports only ADK 10 on Windows Server 2012 to Windows Server 2016. If Windows ADK is not installed on the system, the Create Simplified Disaster Recovery wizard downloads and installs Windows ADK 10.

To create a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2012 to Windows Server 2016

- 1 Do any of the following:
 - On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 2 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 3 Do one of the following:

To create a Simplified Disaster Recovery disk image for the first time	Click Create a new Simplified Disaster Recovery Disk image (.iso) .

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

4 Click **Next**.

5 Do one of the following:

If the ADK is already installed on the computer

If an ADK version other than ADK 10 is installed on the operating system, the Create Simplified Disaster Recovery wizard detects the Windows ADK version and informs the user that a new SDR disk cannot be created using the installed ADK version.

If ADK 10 is installed, the wizard informs you that this is not the latest available version of ADK but you can proceed with the SDR disk creation. Confirm that you want to continue using ADK 10.

After you select the check box and click **Next**, the wizard continues with the recovery disk creation process.

Continue with step [6](#)

If the ADK is not installed on the computer

If Windows ADK is not installed on the system, the Create Simplified Disaster Recovery wizard downloads and installs Windows ADK 10.

Do the following in the order listed:

- Click **Next** to continue to the ADK download and installation wizard.
- On the ADK wizard panels, click **Next** to accept the recommended defaults.

Note: The ADK download may take few hours, depending on the bandwidth.

- After the ADK is installed, click **Close** to continue with the **Create Simplified Disaster Recovery Disk Wizard**.

6 Specify the startup options that you want when you recover a computer, and then click **Next**.

7 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups

Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**.

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include in the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the **Select computers to use the drivers from** panel, click **Next**.

8 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: If you use any non-English characters in the volume label, then those characters are not properly displayed when you mount the .iso image.

Note: It is recommended that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows: `C:\<Backup Exec install path>\Backup Exec\sdr\CustomSDRImage.iso`

9 Review the summary panel, and then click **Create Image**.

- 10 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 11 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Creating a Simplified Disaster Recovery disk image if the Backup Exec server is not available

The Microsoft Assessment and Deployment Kit (ADK) installation is supported only on Windows Server 2012 and later. If the Backup Exec server does not run on Windows Server 2012 or later, or if the server is not available, you can use the Backup Exec installation media to install the **Create Simplified Disaster Recovery Disk Wizard** and the Remote Administration Console on a standalone server that runs Windows Server 2012 or later.

The **Create Simplified Disaster Recovery Disk Wizard** guides you through the download and installation of the ADK. The ADK installation defaults are already selected. It is recommended that you use these defaults. You can then create the recovery disk by connecting to a remote Backup Exec server, as guided by the wizard.

To create a Simplified Disaster Recovery disk image if the Backup Exec server is not available

- 1 Insert the Backup Exec installation media into a server that runs Windows Server 2012 or later.
- 2 From the installation media browser, click **Installation**.
- 3 Click **Simplified Disaster Recovery Disk Creation Wizard**, and then click **Install**.

Note: The **Simplified Disaster Recovery Disk Creation Wizard** and the Backup Exec Remote Administration Console are installed.

- 4 Accept the terms of the license agreement, and then click **Next**.
- 5 Accept the default installation location, or specify a new location, and then click **Next**.
- 6 By default, after the computer is restarted, the **Create Simplified Disaster Recovery Disk Wizard** starts; if you unchecked the default, to start the wizard, click **Start > All Programs > Arctera Backup Exec > Create Disaster Recovery Disk**.
- 7 On the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, click **Connect to a Backup Exec server**.

8 Specify a Backup Exec server machine name or IP address, enter your credentials and the domain, and then click **Next**.

9 Do one of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk Image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

10 Click **Next**.

11 Do one of the following:

If the ADK is already installed on this computer

Continue with step [12](#)

If the ADK is not installed on this computer

Do the following in the order listed:

- Click **Next** to continue to the ADK download and installation wizard.
- On the ADK wizard panels, click **Next** to accept the recommended defaults.

Note: The ADK download may take few hours depending on your network speed.

- After the ADK is installed, click **Close** to continue with the **Create Simplified Disaster Recovery Disk Wizard**.

12 Specify the startup options that you want when you recover a computer, and then click **Next**.

13 Do one of the following:

To add network drivers and storage drivers from the computers for which you ran SDR backups	Select the computers for which you want to add the drivers to the recovery disk, and then click Next
To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image	Do the following in the order listed: ■ On the Select computers to use the drivers from panel, click Next. ■ On the Drivers to Include panel, select the drivers that you want to include on the custom recovery disk, and then click Next.
To add network drivers and storage drivers from OEM media	Do the following in the order listed: ■ On the Select computers to use the drivers from panel, click Next. ■ On the Drivers to Include panel, click Add Driver, navigate to the location of the media, select it, and then click Open. ■ When the drivers display on the Select computers to use the drivers from panel, click Next.
To add a computer to the list	Do the following in the order listed: ■ Click Add DR file. ■ Navigate to the location of the DR file, select it, and then click OK. ■ On the Select computers to use the drivers from screen, select the computer that you added, and then click Next.

14 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: It is recommended that you store the image that you want to burn in a location on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows:

`C:<Backup Exec install path>\Backup Exec\sdr\CustomSDRImage.iso`

- 15 Review the summary panel, and then click **Create Image**.
- 16 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 17 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Creating a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an internet connection and does not have ADK or Windows Preinstallation Environment installed

For Windows ADK 10, you can download the ADK executable file to a computer that has Internet access, and then copy the downloaded files to the computer on which you want to install ADK. Then, you can run the **Simplified Disaster Recovery Disk Creation** wizard to create the recovery disk.

For Windows ADK 1809, you can download the ADK and Windows Preinstallation Environment executable files to a computer that has Internet access, and then copy the downloaded files to the computer on which you want to install the ADK and Windows Preinstallation Environment. Then, you can run the **Simplified Disaster Recovery Disk Creation** wizard to create the recovery disk.

To create a Simplified Disaster Recovery disk image on a Backup Exec server with Windows Server 2019 or Windows Server 2022 installed that does not have an Internet connection and does not have Windows ADK 1809 or Windows Server 2022 compatible ADK and Windows Preinstallation Environment installed

- 1 On a computer that has an Internet connection, for Windows ADK 1809 or Windows Server 2022 compatible ADK, manually download and install Windows ADK using the following link.
<https://docs.microsoft.com/en-us/windows-hardware/get-started/adk-install?ocid=tia-235208000>
- 2 After the download is complete, verify that the adksetup.exe file and the Installers folder have been downloaded to the computer.
- 3 Copy the downloaded files and folder to the computer on which you want to install ADK.
- 4 Start adksetup.exe from the copied folder.
- 5 On the ADK **Specify Location** panel, click **Install the Assessment and Deployment Kit to this computer**, enter the path to which you want to install the ADK, and then click **Next**.
- 6 Click **Accept** to accept the terms of the license agreement.

- 7 On the ADK **Select the features you want to install** panel, click **Deployment Tools**, and then click **Install**.
Complete the ADK installation.
- 8 Verify that the adkwinpesetup.exe file and the Installers folder have been downloaded to the computer.
- 9 Copy the downloaded files and folder to the computer on which you want to install the preinstallation environment.
- 10 Start adkwinpesetup.exe from the copied folder.
- 11 On the ADK **Specify Location** panel, click **Install the Assessment and Deployment Kit Windows Preinstallation Environment Add-ons to this computer**, enter the path to which you want to install the preinstallation environment, and then click **Next**.
- 12 Click **Accept** to accept the terms of the license agreement.
- 13 On the preinstallation environment **Select the features you want to install** panel, select **Windows Preinstallation Environment (Windows PE)**, and then click **Install**.
- 14 After the installation is complete, on the Backup Exec server to which you installed the ADK and preinstallation environment, do any of the following:
 - On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 15 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 16 Do any of the following:

To create a Simplified Disaster Recovery disk image for the first time	Click Create a new Simplified Disaster Recovery Disk image (.iso) .

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

17 Click **Next**.

18 Specify the startup options that you want when you recover a computer, and then click **Next**.

19 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups

Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**.

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include in the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the panel, click **Next**.

- 20 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: It is recommended that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows: `C:<Backup Exec install path>\Backup Exec\sdr\CustomSDRImage.iso`.

- 21 Review the summary panel, and then click **Create Image**.
- 22 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 23 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

To create a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an Internet connection and does not have ADK 10 installed

- 1 On a computer that has an Internet connection, click the following link to download the adksetup.exe:

For Windows ADK 10 <https://go.microsoft.com/fwlink/p/?LinkId=526740> (ADK 10.1.14393.0)
- 2 Run the adksetup.exe.
- 3 On the ADK **Specify Location** panel, click **Download the Assessment and Deployment Kit for installation on a separate computer**, enter the path to the Backup Exec server to which you want to install the ADK components, and then click **Next**.
- 4 Click **Yes** or **No** to join the Microsoft Customer Experience Improvement Program, and then click **Next**.
- 5 Click **Accept** to accept the terms of the license agreement and to start the ADK download.
- 6 After the download is complete, verify that the adksetup.exe file and the Installers folder have been downloaded to the computer.
- 7 Copy the downloaded files and folder to the computer on which you want to install ADK.

Note: ADK installation is supported only on Windows Server 2012 and later.

- 8 Start adksetup.exe from the copied folder.

- 9 On the ADK **Specify Location** panel, click **Install the Assessment and Deployment Kit to this computer**, enter the path to which you want to install the ADK, and then click **Next**.
- 10 Click **Accept** to accept the terms of the license agreement.
- 11 On the ADK **Select the features you want to install** panel, select **Deployment Tools** and **Windows Preinstallation Environment (Windows PE)**, and then click **Install**.
- 12 After the installation is complete, on the Backup Exec server to which you installed the ADK, do any of the following:
 - On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 13 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 14 Do any of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

- 15 Click **Next**.
- 16 Specify the startup options that you want when you recover a computer, and then click **Next**.

17 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups	Select the computers for which you want to add the drivers to the recovery disk, and then click Next .
To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image	Do the following in the order listed: ■ On the Select computers to use the drivers from panel, click Next. ■ On the Drivers to Include panel, select the drivers that you want to include in the custom recovery disk, and then click Next.
To add network drivers and storage drivers from OEM media	Do the following in the order listed: ■ On the Select computers to use the drivers from panel, click Next. ■ On the Drivers to Include panel, click Add Driver, navigate to the location of the media, select it, and then click Open. ■ When the drivers display on the panel, click Next.

18 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: It is recommended that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows: `C:<Backup Exec install path>\Backup Exec\sdr\CustomSDRImage.iso`

- 19 Review the summary panel, and then click **Create Image**.
- 20 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 21 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Contents of the Simplified Disaster Recovery disk image

The Simplified Disaster Recovery disk image provides the **Recover This Computer** wizard that guides you through the steps to recover a computer.

The Simplified Disaster Recovery disk image also includes the tools that you can use to manage the server's network configuration while in the recovery environment. Other utilities are included on the disk that you can use for additional recovery environment operations, such as dynamic loading of drivers, and accessing a command prompt.

Network tools include the following menu items:

- **Start My Networking Services**
- **Map a Network Drive**
- **Configure Network Connection Settings**
- **Run IP Config Utility**
- **Ping a Remote Computer**

Utility tools include the following menu items:

- **Gather Log Files for Technical Support**
- **View Log File**
- **Start PowerShell**
- **Start Command Prompt**
- **Edit the Windows boot.ini File**
- **Load a Driver**
- **Select Keyboard**

Simplified Disaster Recovery also provides you with a hard disk configuration tool called Advanced Disk Configuration. Advanced Disk Configuration lets you run advanced hard disk operations on the computer that you want to recover.

For example, you can do tasks such as:

- Create or delete volumes.
- Resize a volume.
- Convert a basic hard disk to a dynamic disk.
- Change or assign drive letters.
- Create spanned, striped, and mirrored volumes.

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 926.

See [“Requirements for using Simplified Disaster Recovery ”](#) on page 878.

User scenarios when a user starts the Create Recovery Disk Wizard

When a user starts the **Create Simplified Disaster Recovery Disk Wizard**, the wizard first internally checks whether Windows ADK is already installed on the system. Refer the table for a high-level sequence of events.

Table 21-4 User scenarios when a user starts the Create Recovery Disk Wizard

Windows server version	If Windows ADK is not installed	If Windows ADK is installed
Windows Server 2012	If Windows ADK is not installed, the Simplified Disaster Recovery Disk Creation wizard downloads Windows ADK 10 from the Microsoft Web site. After Windows ADK 10 download and install, the wizard starts creating the SDR disk.	If Windows ADK is installed, the Simplified Disaster Recovery Disk Creation wizard verifies if the Windows ADK version is 10. If ADK 10 is installed, the wizard informs the user that this is not the latest ADK version but the user can proceed with SDR disk creation. If the Windows ADK version is below 10, the wizard displays an informational message to the user to download Windows ADK 10 from the Microsoft Web site and then proceed with the SDR disk creation. Wizard does not support creating a new disk using ADK 8.0 or ADK 8.1.
Windows Server 2012 R2	Same as Windows Server 2012	Same as Windows Server 2012
Windows Server 2016	Same as Windows Server 2012	Same as Windows Server 2012

Table 21-4

User scenarios when a user starts the Create Recovery Disk Wizard (continued)

Windows server version	If Windows ADK is not installed	If Windows ADK is installed
Windows Server 2019	If Windows ADK 1809 is not installed, the user must download and install Windows ADK 1809 manually. When you install ADK 1809, you must install the Windows ADK Insider and Windows Preinstallation Environment components. After the components are installed, launch the Simplified Disaster Recovery Disk Creation wizard again.	If Windows ADK 1809 is installed, the Simplified Disaster Recovery Disk Creation wizard detects the installed version of ADK and the wizard continues with the recovery disk creation process. If Windows ADK 10 is installed, the Simplified Disaster Recovery Disk Creation wizard detects that ADK 10.0 version is installed and the ADK update screen is displayed. The wizard informs the user that this is not the latest available version of ADK but the user can proceed with the SDR disk creation. Confirm that the user wants to continue using ADK 10. If the user continues, it may not support all the features of Windows Server 2019.

Table 21-4

User scenarios when a user starts the Create Recovery Disk Wizard (continued)

Windows server version	If Windows ADK is not installed	If Windows ADK is installed
Windows Server 2022	If Windows ADK 1809 is not installed, the user must download and install the Windows ADK manually. When you install the Windows ADK version, you must install the Windows ADK Insider and Windows Preinstallation Environment components. After the components are installed, launch the Simplified Disaster Recovery Disk Creation wizard again.	If Windows ADK 1809 is installed, the Simplified Disaster Recovery Disk Creation wizard detects the installed version of ADK and the wizard continues with the recovery disk creation process. If Windows ADK 10 is installed, the Simplified Disaster Recovery Disk Creation wizard detects that ADK 10.0 version is installed and the ADK update screen is displayed. The wizard informs the user that this is not the latest available version of ADK, but the user can proceed with the SDR disk creation. Confirm that the user wants to continue using ADK 10. If the user continues, it may not support all the features of Windows Server 2019 or Windows Server 2022.

Table 21-4

User scenarios when a user starts the Create Recovery Disk Wizard *(continued)*

Windows server version	If Windows ADK is not installed	If Windows ADK is installed
Windows Server 2025	If Windows ADK 1809 is not installed, the user must download and install the Windows ADK manually. When you install the Windows ADK version, you must install the Windows ADK Insider and Windows Preinstallation Environment components. After the components are installed, launch the Simplified Disaster Recovery Disk Creation wizard again.	If Windows ADK 1809 is installed, the Simplified Disaster Recovery Disk Creation wizard detects the installed version of ADK and the wizard continues with the recovery disk creation process. If Windows ADK 10 is installed, the Simplified Disaster Recovery Disk Creation wizard detects that ADK 10.0 version is installed and the ADK update screen is displayed. The wizard informs the user that this is not the latest available version of ADK, but the user can proceed with the SDR disk creation. Confirm that the user wants to continue using ADK 10. If the user continues, it may not support all the features of Windows Server 2019, Windows Server 2022, or Windows Server 2025.

Preparing to recover from a disaster by using Simplified Disaster Recovery

When a disaster occurs, you can use Simplified Disaster Recovery (SDR) to return the computer to its pre-disaster state.

To prepare to recover a computer, use the following steps:

Table 21-5

Preparing to recover from a disaster by using Simplified Disaster Recovery

Step	Description
Step 1	Disconnect any storage area network or cluster that is attached to the computer being recovered. Otherwise, the hard drives on those computers may also be repartitioned and reformatted.
Step 2	Plan any hardware changes to the computer to be recovered. See “Hardware replacement during disaster recovery” on page 914.
Step 3	Review additional requirements for IBM computers if the computer to be recovered is an IBM computer. See “Prepare to recover IBM computers with Simplified Disaster Recovery” on page 915.
Step 4	Review recovery notes if you are recovering a SQL server, Exchange server, SharePoint server, or an environment in which the Central Admin Server feature is installed. See “Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature” on page 924. Review recovery notes if you are recovering Windows operating systems with storage pools and storage spaces. See “Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces” on page 922.

When you complete these steps, you can start the recovery process.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 916.

Hardware replacement during disaster recovery

You can use Simplified Disaster Recovery (SDR) to recover a computer that is no longer functioning. For example, if the computer’s main system board fails, you

can restore the computer’s data after you replace the system board. You can also restore the data even if the new board is a different model or if it contains multiple processors.

If you plan to change the hardware in the computer being recovered, consider the following information before you use SDR to recover the computer:

Table 21-6 Hardware considerations when recovering failed computers

Item	Description
Hard drives	Any hard drives that you replace should be the same size or larger than the original drives.
System boards	After you replace a faulty system board and after you use SDR to recover the computer, you must use the system board manufacturer’s driver CD to re-install additional functionality such as onboard sound and video.
Network interface cards	If you change the network interface card in the computer that you want to recover, you must install the necessary network drivers. Without the network drivers, you cannot access the network if you want to use a remote Backup Exec server or remote legacy backup-to-disk folders to recover the computer. After you complete the recovery, you must install new network interface card drivers that match the network card that is currently in the computer. In most instances, it is not necessary to install the drivers manually. Most drivers are available on the Simplified Disaster Recovery disk image. When you create the Simplified Disaster Recovery disk image, you can customize it to include any drivers that you may need.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 916.

Prepare to recover IBM computers with Simplified Disaster Recovery

To prepare to recover an IBM computer that has an IBM ServerAID card by using Simplified Disaster Recovery, use the following steps:

Table 21-7 Prepare to recover an IBM computer

Step	Description
Step 1	Install and configure the IBM ServeRAID controller card and ServeRAID software so that a boot volume is recognizable by the Windows operating system.
Step 2	Before you use the SDR startup media, start the server by using the IBM server's ServeRAID Configuration and Management CD in the CD-ROM drive. This starts the IBM ServeRAID utilities configuration and installation process to view and update the current BIOS and firmware levels.

Refer to the IBM ServeRAID documentation for complete installation instructions for installing Windows on an IBM Server with the ServeRAID controller. Create and initialize the ServeRAID disks so that volumes are recognizable by the Windows operating system.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 916.

Recovering a computer with Simplified Disaster Recovery

You can use the Simplified Disaster Recovery **Recover This Computer Wizard** to run a local recovery or a remote recovery of a Backup Exec server or a Windows computer.

You can perform a local recovery if all of the following conditions are met:

- You want to restore a Backup Exec server or a Windows computer.
- The backup data for the computer is located on the devices that you can locally attach to the computer on which you run the **Recover This Computer Wizard**.
- You can provide the disaster recovery information file for the Backup Exec server or the Windows computer that you want to recover.

Note: If the computer that you want to recover was backed up to a tape device, to a deduplication storage, or to a virtual disk, then SDR was unable to store the disaster recovery information file with the backup sets. You must provide the path to the default location or to the alternate location for the disaster recovery information file when the **Recover This Computer Wizard** prompts you. If you cannot provide a disaster recovery information file, then you must perform a manual disaster recovery. Refer to the following sections for more information about how to perform manual disaster recovery:

See [“Performing manual disaster recovery of a local Backup Exec server on a Windows computer”](#) on page 930.

See [“Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer”](#) on page 935.

Note: If you are restoring data on the same computer and if you do not select the **Erase hard disks and recreate the volume layout shown above** option, the recovery process does not restore the WindowsApps folders on operating systems that run Windows 8 or later. However, the restore job is successful. Microsoft recommends to restore the WindowsApps folders using Device Reset on the Settings panel.

The WindowsApps folders that are ignored during restore could be the following:

The folder pointed by

`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Appx\PackageRoot`

The folder pointed by

`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Appx\PackageRepositoryRoot`

`%SystemRoot%\InfusedApps`

You can perform a remote recovery if all of the following conditions are met:

- You want to restore a Backup Exec server or a Windows computer.
- The backup data for the computer is located on the devices that are attached to a remote Backup Exec server
- You have network access to the remote Backup Exec server.

Before you start, review the steps for preparing to recover.

See [“Preparing to recover from a disaster by using Simplified Disaster Recovery”](#) on page 913.

During recovery operations, you can recover the computer from the most recent backup. Or, you can recover to a previous point-in-time backup.

The **Recover This Computer Wizard** supports the recovery of computers with encrypted backup sets. If the Simplified Disaster Recovery (SDR) backups are encrypted during backup, the wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the recovery.

See [“Encryption key management”](#) on page 733.

To restore data by using the **Recover This Computer Wizard**, the following items are required:

- The Simplified Disaster Recovery disk image, which must be the same version of Backup Exec as the Backup Exec server.
- If the computer that you want to recover was backed up to a tape device, deduplication storage, or to a virtual disk, then the disaster recovery information file is not stored with the backup sets. Instead, you must provide the path to the default location or to the alternate location when the **Recover This Computer Wizard** prompts you. The default location is: C:<Backup Exec install path>\Backup Exec\sdr\Data. If the file is unavailable, you cannot recover the computer with SDR.
- The backup set that contains all of the critical system components for the computer that you want to restore.
- The boot drive on the computer that you want to recover must have from 3 GB to 5 GB of free space, depending on the operating system and configuration.

Note: If a blank screen appears and the computer does not restart after you run SDR, ensure that the boot drive has the necessary amount of free space. Then, restart the computer again.

- The computer that you want to recover must have an amount of RAM that is equal to or greater than the minimum that the restored operating system requires.
- The computer that you want to recover must have disks with enough free space to contain the data that you restore from the backup of critical volumes. A volume is considered critical if it is required for the computer to start successfully.
- If you recover a Windows computer that has BitLocker encryption enabled, you must turn on BitLocker encryption after the restore. See your Microsoft documentation for more information on BitLocker drive encryption.

- If the computer that you want to recover contains a RAID setup, you may be required to configure the RAID before you start it with the SDR disk. Use the computer manufacturer's RAID software to configure the RAID system.
- If appropriate, review the following recovery notes:
 - See [“Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces”](#) on page 922.
 - See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature”](#) on page 924.

Note: Boot managers, such as System Commander or the OS/2 Boot Manager, cannot be restored with SDR. Boot managers are usually installed at a very low level that Backup Exec cannot protect. For example, the OS/2 Boot Manager resides in its own hard drive volume that Backup Exec cannot access. Because of the many different boot managers available, you may not be able to restart the computer after an SDR recovery, even though the operating system was restored. You can reinstall the boot manager to resolve this issue.

To recover a computer by using the Recover This Computer Wizard

- 1 Place the startup Simplified Disaster Recovery Disk in the CD drive or DVD drive of the computer that you want to recover, and then start the computer.
- 2 On the **End User License Agreement** screen, click **Accept**.
- 3 On the Simplified Disaster Recovery **Welcome** screen, click **Recover This Computer**.
- 4 In the left pane, click **Network**, and then click **Configure Network Connection Settings**.
- 5 Select the appropriate network adapter configuration, and then click **OK**.

Note: If required, set the correct local time zone and local time using Windows PowerShell.

- 6 Do one of the following:

If the backup data for this computer is located on the devices that are attached to a remote Backup Exec server (remote recovery)

Do the following in the order listed:

- Click **The data is located on devices attached to a remote Backup Exec server.**
- Enter the name and domain of the remote Backup Exec server where the backup data is located, and the appropriate administrator or administrator-equivalent credentials.
- (Optional) To configure network adapter settings, click **Configure network adapter settings**, and then do one of the following:
 - To assign a static IP address for each detected network adapter, select the appropriate options, and then click **OK**.
 - To configure an IPv6 network controller, click **Configure IPv6**, select the appropriate options, and then click **OK**.
- (Optional) To load network adapter drivers, click **Load network adapter drivers**. Then, click **Install Driver** for any inactive network controller that the wizard detects. Navigate to the device that contains the network controller driver, and then click **Open**. Select the driver, and then click **Open** again.
- Click **Next**.

If the computer that you want to recover was backed up to locally attached devices such as tape drives, robotic libraries, disk storage devices, or disk cartridge devices (local recovery)

Do the following in the order listed:

- Click **The data is located on devices locally attached to this computer.**
- Enter the appropriate administrator or administrator-equivalent credentials when prompted, and then click **Next**.

- 7 Select the disaster recovery information file that you want to use, and then click **Next**.

Note: Backup Exec by default supports the latest three full SDR backup chains that are required for system recovery using SDR. Each backup chain includes one full backup set, its dependent incremental and differential backups, and their duplicate backup sets.

- 8 Select the backup sets that you want to use to recover the computer, and then click **Next**.
- 9 Select the storage device that contains the backup data for the computer, and then click **Next**.
- 10 Do any of the following:

To use the volume layout as shown Click **Next**.

To view the disk geometry as it presently exists, or to view a graphical representation of proposed changes Click **Preview**.

To install required SCSI or RAID controller drivers

Do the following in the order listed:

- Click **Load Storage controller drivers**.
- Click **Install Driver** for any inactive network controller that the wizard detects.
- Navigate to the device that contains the network controller driver, and then click **Open**.
- Select the driver, and then click **Open** again.

To automatically create a volume layout on the available hard disks if mismatched volumes appear in the simplified volume layout view Select **Erase hard disks and recreate the volume layout shown above**. Existing data on these disks will be lost.

To change volume sizes or other disk-related operations

Click **Advanced Disk Configuration**.

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 926.

- 11 Review the **Recovery Summary** and click **Back** to make any changes, or click **Recovery** to start the recovery process.
- 12 If you recover a Backup Exec server, select the appropriate database encryption key and then click **Next**.

Note: Backup Exec requires a database encryption key to access the Backup Exec Database. If you proceed without entering the database encryption key on this screen, you are prompted to enter the key when you restart Backup Exec.

- 13 When the recovery is complete, do any or all of the following:

To troubleshoot any issues that occurred during the hardware discovery phase Click **View Hardware Discovery log**.

To troubleshoot any issues that occurred during the recovery process Click **View Recover This Computer log**.

To restart this computer after you click **Finish** Select **Restart this computer**.

To complete the recovery process and close the wizard Click **Finish**.

Note: If the recovered computer contains multiple hard disks, ensure that the computer's BIOS is configured to start the computer from the hard disk that contains the Windows operating system. On many computers, the BIOS menu can be accessed by pressing F2 or DEL on the computer keyboard.

Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces

You can perform Simplified Disaster Recovery (SDR) for Windows operating systems with storage pools and storage spaces using either of the scenarios that are described in the following table:

Table 21-8 Scenarios for SDR for storage pools and storage spaces

Scenario	Description
Restore Windows operating systems with storage pools and storage spaces to the same computer	In this scenario, the Recover This Computer Wizard prompts you to reconfigure storage pools and storage spaces, or remove them, or to restore them as they are. To reconfigure or remove the storage pools and storage spaces, click the PowerShell button on the displayed message dialog box to open a PowerShell window. If the Backup Exec Storage folder (BEData) is configured on the storage pools and storage spaces, you may not be able to do an SDR Local recovery from the BEData folder. This can happen when the installed ADK version and the operating system version of the Backup Exec server do not match. Storage pools and spaces may not be detected during system restore if the recovery disk was created using an ADK version that was not released along with the operating system version being recovered. For more information on ADK version and corresponding operating system, refer to the Microsoft documentation. For more information about storage commands in PowerShell, see the following URL: http://technet.microsoft.com/library/hh848705.aspx

Limitations for restoring storage pools and storage spaces

The **Recover This Computer Wizard** automatically maps volumes to their original virtual disks or storage spaces if you restore to the same computer. If you do not restore to the same computer, then you must use PowerShell and the SDR Advanced Disk Configuration utility to manually map the volumes to the virtual disks or storage spaces.

The **Recover This Computer Wizard** does not allow Boot, System, and Recovery volumes to be mapped to virtual disks.

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 926.

Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature

Review the following recovery notes if you recover a SQL Server, Exchange Server, SharePoint server, or an environment in which the Central Admin Server feature is installed.

Microsoft SQL Server recovery notes

After you use Simplified Disaster Recovery (SDR) to recover the Windows server, SDR automatically replaces the damaged master and model databases with copies of the master database and the model database. After SQL is restarted and the latest master database backup and all other system databases are restored, you must still restore all user databases. Use the Backup Exec Restore wizard to restore the latest backups.

Microsoft Exchange recovery notes

After you use SDR to recover the Windows server, use the Backup Exec Restore wizard to restore the Exchange Server databases from the most recent Exchange Server database backups.

SharePoint Portal Server recovery notes

You can use SDR to recover a Windows server that has SharePoint Portal Server installed. After you restore the Windows computer, you must restart it. After the computer restarts, although, the SharePoint Portal Server software is installed, it is not functional. You must remove the SharePoint Portal Server and reinstall it before the SharePoint data can be restored.

Central Admin Server feature (CAS) recovery notes

When you use SDR to recover a computer in a CAS environment, you can submit the remote restore job to either of the following:

- The central administration server.
- The managed Backup Exec server that performed the original backup job.

Microsoft Hyper-V hosts

After you use SDR to recover the Windows server, use the Backup Exec Restore wizard to restore the Microsoft Hyper-V guest virtual machines from the most recent Microsoft Hyper-V backups.

Backup Exec Deduplication feature recovery notes

If the Backup Exec Deduplication feature is installed and you use deduplication disk storage devices, consider the following:

- If you back up the local Backup Exec server, do not select a deduplication disk storage device on the server as the destination storage device. The **Recover This Computer** wizard cannot restore data from a local deduplication disk storage device.
- Simplified Disaster Recovery (SDR) cannot recover a deduplication disk storage device.
- Before you can use SDR to restore a remote computer that was backed up with client-side deduplication, you must first delete the direct access device. See [“Selecting storage devices for direct access sharing”](#) on page 991.

If you use SDR to recover a Backup Exec server that contains a deduplication disk storage device, consider the following:

- Any existing backup sets that were sent to the deduplication disk storage device after it was backed up cannot be restored.
- The deduplication disk storage device may not be in an operational state after the recovery.

You can also use manual disaster recovery to recover deduplication disk storage devices and OpenStorage devices.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“Disaster recovery of deduplication disk storage devices”](#) on page 1001.

See [“Disaster recovery information file data paths”](#) on page 890.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 889.

See [“About the Central Admin Server feature”](#) on page 1321.

Recovery notes for using Simplified Disaster Recovery with Windows BIOS system

Review the following recovery note if you want to perform Simplified Disaster Recovery (SDR) for Windows BIOS system.

- With Windows Server 2022 BIOS operating systems, a drive letter is assigned to a recovery partition after successful restore using Simplified Disaster Recovery. You can either remove the recovery partition before starting restore using SDR or you can manually remove the drive letter after the system boots up.

Advanced Disk Configuration on the Recover This Computer Wizard

The **Recover This Computer Wizard** restores the hard drive volumes to the same sizes they were before the disaster. If the hard drive in the failed computer is larger than the hard drive that was in place before the disaster, it may result in unused and unallocated space. You can run Advanced Disk Configuration to alter the volume sizes to reflect the larger hard drive size.

The following is an example of why the hard drive volumes should be resized:

The pre-disaster computer hardware contains a 40 GB hard drive with two 20-GB volumes. You replace it with a 90-GB model. SDR then uses the disaster recovery information file to rebuild the hard disk partition table by using the partition information that is found on the original 40-GB hard drive. As a result, only 40 GB of space is allocated on the new 90 GB hard drive, with a partition map that consists of two 20-GB partitions.

You can access Advanced Disk Configuration from within the **Recover This Computer Wizard**.

Note: You should be familiar with Microsoft Disk Management concepts before you run Advanced Disk Configuration.

The following table provides details about the additional disk-related operations that you can do with Advanced Disk Configuration.

Table 21-9 Advanced Disk Configuration tasks

Task	Description
Create a simple volume	A simple volume is a partition on a disk that contains a file system.
Format a volume	Disk volumes must be formatted before data can be stored on them.
Extend the size of a volume	If a disk contains some unallocated disk space that is adjacent to a functional volume, you can extend the volume to include the free space. To extend the volume, it must be either raw or formatted with the Windows NTFS file system.

Table 21-9 Advanced Disk Configuration tasks (*continued*)

Task	Description
Shrink the size of a volume	You can decrease the size of a volume by shrinking the volume into the contiguous, unallocated disk space that is on the same disk. When you shrink a volume, there is no need to reformat the volume. Ordinary files are automatically relocated on the disk to create the new, unallocated disk space.
Create a spanned volume	A spanned volume spans more than one physical disk. You can create a spanned volume by spanning it across multiple physical disks, or by spanning the volume into unallocated disk space. To create a spanned volume, you must have a startup volume and at least two dynamic volumes. Note: Spanned volumes are not fault-tolerant.
Create a striped volume	Striped volumes store data in stripes across two or more physical disks. Although striped volumes do not provide fault-tolerance protection, they do offer the best performance of all the volumes in Windows.
Create a mirrored volume	A mirrored volume provides a copy of the data that is written to a selected volume. Because all data is written to both the mirrored volume and the selected volume, mirroring reduces the capacity of both volumes by 50%.
View volume properties	You can view properties for each volume in the Current Disk Layout view or in the Original Disk Layout view.
Change an assigned drive letter	You can change assigned drive letters for all volumes if you want to organize your drive letters in a certain way.

Table 21-9 Advanced Disk Configuration tasks (continued)

Task	Description
Delete a volume	Deleting a volume erases all data from the volume; therefore, it is recommended that you use caution when considering the use of this option.
Convert a basic disk to a dynamic disk	Converting basic disks to dynamic disks lets you create the volumes that span multiple disks. Dynamic disks also let you create fault-tolerant volumes, such as mirrored volumes and RAID-5 volumes. All volumes on dynamic disks are referred to as dynamic volumes.
Convert a Master Boot Record (MBR) disk to a Guid Partition Table (GPT) disk	MBR disks use the standard BIOS interface. GPT disks use extensible firmware interface (EFI). You can convert MBR disks to GPT disks if the disk does not contain partitions or volumes.
Convert a Guid Partition Table (GPT) disk to a Master Boot Record (MBR) disk	GPT disks use extensible firmware interface (EFI). Master boot record (MBR) disks use the standard BIOS interface. GPT disks can be converted to MBR disks if the disk does not contain partitions or volumes.

Table 21-9 Advanced Disk Configuration tasks (*continued*)

Task	Description
View the original disk layout geometry	The original disk layout shows the actual hard disk layout that existed during the backup job. Within the simplified layout view, you can accept the disk geometry as it originally existed before the disaster, or you can alter the geometry by changing the volume sizes. Depending on the size of the existing disks, you can alter volume sizes in megabytes, gigabytes, or terabytes. The simplified layout view has a Preview tab that lets you view the disk geometry as it presently exists. If you change the disk geometry and change volume sizes, click the Preview tab to see a graphical representation of your proposed changes. If mismatched volumes appear in the simplified volume layout view, you can use the option Erase hard disks and recreate the volume layout shown above to automatically create a volume layout on the available hard disks. You can also manually create a volume layout by using the Advanced Disk Configuration option.

Performing manual disaster recovery

You should perform manual disaster recovery in the following situations:

- Simplified Disaster Recovery is not available or has failed.
- The Windows operating system has become corrupted and cannot be repaired using the Emergency Repair Disks.
- The hard drive containing the Windows operating system has encountered an unrecoverable error that requires reformatting the disk.
- The hard drive that contains the Windows operating system needs to be replaced.
- The server has been hit by ransomware or any other virus attacks and it is better to recover the server from a good recent backup.

Note: For Windows 2025, use Simplified Disaster Recovery (SDR) for faster and reliable recovery experience. Manual DR is available, but requires user steps after recovery. For more details, see article [100074698](https://www.veritas.com/support/en_US/article.100074698) (https://www.veritas.com/support/en_US/article.100074698).

For information about how to perform manual disaster recovery of a local Backup Exec server on a Windows computer, see the following section. This can be a member server or a domain controller.

See [“Performing manual disaster recovery of a local Backup Exec server on a Windows computer”](#) on page 930.

For information about how to perform manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer, see the following section. This can be a member server or a domain controller.

See [“Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer”](#) on page 935.

Performing manual disaster recovery of a local Backup Exec server on a Windows computer

This procedure restores the computer's operating system to a pre-disaster state. It also restores the data files, except for those that a Backup Exec agent protects, such as the Agent for Microsoft Exchange Server, the Agent for Microsoft SQL Server, or the Agent for Microsoft Hyper-V . If Backup Exec agents protect any of the data, refer to the section in the Backup Exec Administrator's Guide on how to restore the data that is protected by the agent before beginning disaster recovery. The agent-protected data should be restored after the system recovery is complete. This procedure includes non-authoritative and authoritative restore of a domain controller.

These steps are intended for manual disaster recovery only. If Simplified Disaster Recovery (SDR) is enabled for the computer, you should use SDR for disaster recovery.

The following items are required for manual disaster recovery of a local system:

- The latest full backup of the critical components of the server to be recovered and any subsequent incremental and differential backups to ensure server is recovered to its latest backup set state.
- The Windows installation media to install the same version of the operating system that existed at the time of the backup.
- The Backup Exec installation media to install Backup Exec for reading from the backed-up media and performing the recovery.

- The database encryption key that was used to encrypt the Backup Exec Database. You should have exported the key to a secure location. You must retrieve it from that location to complete the recovery process.
- A storage device such as a tape drive, disk storage device, or a robotic library must be attached to the computer that you want to recover.
- If you are recovering a domain controller, irrespective of whether an authoritative restore or non-authoritative restore is being performed, DSRM credentials are needed as the recovery will be performed by booting the freshly installed server into Directory Service Restore Mode. If recovery is performed in normal mode, the domain controller may not boot post the recovery. The Active Directory Domain services must be installed from the **Add Roles and Features** wizard (Server Roles section). **Promote this server to a domain controller** is not required.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must enable BitLocker encryption again after the restore. See Microsoft's documentation for more information on BitLocker drive encryption.

Always log on to Windows using the Administrator account or its equivalent during this procedure.

To run a manual disaster recovery of a local Backup Exec server on a Windows computer

1. Install the original version of Windows, which existed at the time of backup. The same Service Pack and patches need to be applied after Windows is installed.

Note the following scenarios:

- If you recover from an entire hard disk failure, use Windows setup to partition and format the new disk during installation. This Windows installation is necessary to provide Backup Exec with a target to which it can restore the system. The computer name, Windows directory, and the file system (such as NTFS) must be the same as the previous Windows installation. This basic installation is overwritten by the backed-up version, which restores your system configuration, application settings, and security settings.
- If the system was a domain controller, in a specific domain or workgroup, do not join the domain or workgroup. Use the **More...** option on the **Computer Name Change** dialog box to manually add a domain suffix to the computer name that matches the system's original domain or workgroup suffix.

Change the new system name to match the original system name by performing the following steps in the order listed:

- From **System Properties**, on the **Computer Name** tab, click **Change**.

Note: If you are recovering a domain controller, irrespective of whether an authoritative restore or non-authoritative restore is being performed, DSRM credentials are needed as the recovery will be performed by booting the freshly installed server into Directory Service Restore Mode. If recovery is performed in normal mode, the domain controller may not boot post the recovery. The Active Directory Domain services must be installed from the **Add Roles and Features** wizard (Server Roles section). **Promote this server to a domain controller** is not required.

- On the **Computer Name/Domain Changes** dialog box, click **More**.
 - If necessary, select **Change primary DNS suffix when domain membership changes**, and then click **OK**.
 - Restart the system.
2. Install Backup Exec to a directory other than where it was originally installed (a temporary installation). Always log on to Windows using the Administrator account or its equivalent during this procedure.

Note: After recovery is complete, this temporary installation of Backup Exec can be renamed as registry, services, and installation directory are recovered from the backup set.

3. Start Backup Exec, and then add the required Storage device by selecting the **Storage** tab and then **Configure Storage**.

This storage device is the tape where your backup set resides or the disk path where your disk storage device backup files are located.

Note: If you use a disk storage device to recover the local Backup Exec server, do not include the original disk storage device. If you cannot avoid restoring it, you need to ensure that the disk storage device being used for the recovery does not conflict with the original disk storage device location.

4. On the **Storage** tab, click **Inventory and Catalog** to both inventory and catalog the media that contains the latest full, incremental, and differential backups of the computer that you want to recover.
5. Select the **Backup and Restore** tab, and then click **Restore**.
6. Do one of the following:
 - If the restore method **Complete online restore of a computer, or restore system component** is available, do the following in the order listed:
 - Click **Complete online restore of a computer, or restore system component**, and then click **Next**.
 - Click **A Microsoft Windows computer that was fully selected for backup**, and then click **Next**.
 - Select the backup sets that you want to restore, and then click **Next**.
 - Make sure that you deselect the location for restore where your disk storage device backup files were located, or else the restore operation will overwrite them. Applications and data drives can be restored later once the server recovery is completed.
 - Ensure that the option, **Restore over existing files** is selected, and then accept the default selections on the **How do you want to maintain file integrity, hierarchy, and security for restore data** panel.
 - Click **Next**.
 - On the **How do you want to restore operating system features** panel, click **Next**.
 - For an authoritative restore of a domain controller, select the **Mark this server as the primary arbitrator for replication when restoring SYSVOL in the System State** option on the **How do you want to restore System State data?** panel.
 - On the **What additional tasks do you want to perform before and/or after a restore** panel, select any additional tasks that you want to run before or after a restore, and then click **Next**.
 - Schedule the job to run, and then click **Next**.
 - On the **Restore** summary panel, click **Finish**.
 - Do not restart the computer after the restore job finishes.
 - If the restore method **Complete online restore of a computer, or restore system component** is not available, create a restore job and manually select individual system components for recovery. Do not restart the computer after the restore job finishes.

7. Your computer's operating system is now restored to a pre-disaster state, but you should not restart your system yet. Your data files have been restored if they were included in a restore job, except those protected by Backup Exec database agents.

Continue with one of the following:

- For an authoritative restore of a domain controller, skip to step 8.
 - If you want to restore a standalone server or a non-authoritative restore of a domain controller, the recovery is complete. Restart the computer after the restore job successfully completes. If you have copied disk storage device files to another location for the purpose of a restore, you can remove them. Skip to step 9 to complete this procedure.
8. For an authoritative restore of a domain controller, do the following:

Note: Make sure that the system is started into Directory Services Restore Mode for the first restart after the restore. Failing to do so may replicate the Active Directory once the Active Directory services are online. To prevent this replication, you can isolate the system from the network temporarily.

- Press **F8** during startup. A menu appears that lets you diagnose and fix system startup problems.
- Select **Directory Services Restore Mode**.
- Log on using your DSRM credentials.
- Open a command prompt.
- Type `NTDSUTIL`, and then press **Enter**. For more information about running NTDSUTIL for Windows Server, see Microsoft's documentation.
- Type `Activate Instance NTDS`, and then press **Enter**.
- Type `Authoritative Restore`, and then press **Enter**.
- Type the following command, and then press **Enter**:
`restore subtree ou=OU_Name,dc=Domain_Name,dc=xxx`
In this command, `<ou_name>` is the name of the organizational unit that you want to restore, `<domain_name>` is the domain name that the OU resides in, and `<xxx>` is the top-level domain name of the domain controller, such as `com`, `org`, or `net`.
- Repeat these steps as many times as necessary for the specific objects that you need to restore.

- After you have finished restoring Active Directory information, exit NTDSUTIL.
- Restart the computer.

Note: If you have copied disk storage device files to some other location to restore them, they can be removed.

9. When you launch Backup Exec, the program prompts you for the database encryption key file. Complete the following steps to import the database encryption key file:
 - Locate the database encryption key from the secure location to which you backed it up. Backup Exec indicates the name of the key that needs to be restored.
 - Copy the file and then paste it in the Data folder in the directory in which you installed Backup Exec.
 - Log on to Backup Exec.

Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer

This procedure restores the computer's operating system to a pre-disaster state. It also restores the data files, except for those that a Backup Exec agent protects, such as the Agent for Microsoft Exchange Server, the Agent for Microsoft SQL Server, or the Agent for Microsoft Hyper-V. If Backup Exec agents protect any of the data, refer to the section in the Backup Exec Administrator's Guide on how to restore the data that is protected by the agent before beginning disaster recovery. The agent-protected data should be restored after the system recovery is complete. This procedure includes non-authoritative and authoritative restore of a domain controller.

These steps are intended for manual disaster recovery only. If Simplified Disaster Recovery (SDR) is enabled for the computer, you should use SDR for disaster recovery.

The following items are required for manual disaster recovery of a remote system:

- The latest full backup of the critical components of the server to be recovered and any subsequent incremental and differential backups to ensure server is recovered to its latest backup set state.
- The Windows installation media to install the same version of the operating system that existed at the time of the backup.

- The Backup Exec installation media to install Backup Exec for reading from the backed-up media and performing the recovery.
- The database encryption key that was used to encrypt the Backup Exec Database. You should have exported the key to a secure location. You must retrieve it from that location to complete the recovery process.
- If you are recovering a domain controller, irrespective of whether an authoritative restore or non-authoritative restore is being performed, DSRM credentials are needed as the recovery will be performed by booting the freshly installed server into Directory Service Restore Mode. If recovery is performed in normal mode, the domain controller may not boot post the recovery. The Active Directory Domain services must be installed from the **Add Roles and Features** wizard (Server Roles section). **Promote this server to a domain controller** is not required.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must enable BitLocker encryption again following the restore. See Microsoft's documentation for more information on BitLocker drive encryption.

Always log on to Windows using the Administrator account or its equivalent during this procedure.

To run a manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer

1. At the remote computer, install the original version of Windows, which existed at the time of backup. The same Service Pack and patches need to be applied after Windows is installed. Note the following scenarios:
 - If you recover from an entire hard disk failure, use Windows setup to partition and format the new disk during installation. This Windows installation is necessary to provide Backup Exec with a target to which it can restore the system. The computer name, Windows directory, and the file system (such as NTFS) must be the same as the previous Windows installation. This basic installation is overwritten by the backed-up version, which restores your system configuration, application settings, and security settings.
 - If the system was a domain controller, in a specific domain or workgroup, do not join the domain or workgroup. Use the **More...** option on the **Computer Name Change** dialog box to manually add a domain suffix to the computer name that matches the system's original domain or workgroup suffix.

Change the new system name to match the original system name by performing the following steps in the order listed:

- From **System Properties**, on the **Computer Name** tab, click **Change**.

Note: If you are recovering a domain controller, irrespective of whether an authoritative restore or non-authoritative restore is being performed, DSRM credentials are needed as the recovery will be performed by booting the freshly installed server into Directory Service Restore Mode. If recovery is performed in normal mode, the domain controller may not boot post the recovery. The Active Directory Domain services must be installed from the **Add Roles and Features** wizard (Server Roles section). **Promote this server to a domain controller** is not required.

- On the **Computer Name/Domain Changes** dialog box, click **More**.
 - If necessary, select **Change primary DNS suffix when domain membership changes**, and then click **OK**.
 - Restart the system.
2. Perform a local install of Backup Exec Remote Agent for Windows on the remote computer. Choose a different drive or folder other than where agent was installed earlier. Normally it is `C:\Program Files\Veritas\Backup Exec\RAWS\`, which the default path.

Note: After recovery, the Backup Exec logon account needs to be updated and the Backup Exec trust needs to be reestablished for the recovered remote server.

3. On the **Backup and Restore** tab, select the computer name, and then click **Restore**.
4. Do one of the following:

If the restore method **Complete online restore of a computer, or restore system component** is available, do the following in the order listed:

- Click **Complete online restore of a computer, or restore system component**, and then click **Next**.
- Click **A Microsoft Windows computer that was fully selected for a backup**, and then click **Next**.

- Select the point in time, ensuring that only the critical set is selected, and then click **Next**.
 - Select **To the original location**, and then click **Next**.
 - Ensure that the option, **Restore over existing files** is selected, and then accept the default selections on the **How do you want to maintain file integrity, hierarchy, and security for restore data** panel.
 - Click **Next**.
 - For an authoritative restore of a domain controller, select **Mark this server as the primary arbitrator for replication when restoring SYSVOL in the System State** option on the **How do you want to restore System State data?** panel.
 - On the **What additional tasks do you want to perform before and/or after a restore?** panel, select any additional tasks that you want to run before or after a restore, and then click **Next**.
 - Schedule the job to run and then click **Next**.
 - On the **Restore** summary panel, click **Finish**.
 - Do not restart the computer.
If the restore method **Complete online restore of a computer, or restore system component** is not available, create a restore job and manually select individual system components for recovery. Do not restart the computer.
5. Your computer's operating system is now restored to a pre-disaster state, but you should not restart your system yet. Your data files have been restored if they were included in a restore job, except those protected by Backup Exec database agents.
- Continue with one of the following:
- For an authoritative restore of a domain controller, proceed to step 6.
 - If you want to perform a standalone server restore or a non-authoritative restore of a domain controller, the recovery is complete. Restart the computer after the restore job successfully completes. Proceed to step 7 to complete this procedure.
6. For an authoritative restore of a domain controller, do the following:

Note: Make sure that the system is started into Directory Services Restore Mode for the first restart after the restore. Failing to do so may replicate the Active Directory once the Active Directory services are online. To prevent this replication, you can isolate the system from the network temporarily.

- Press **F8** during startup. A menu appears that lets you diagnose and fix system startup problems.
 - Select **Directory Services Restore Mode**.
 - Log on using your DSRM credentials.
 - Open a command prompt.
 - Type `NTDSUTIL`, and then press **Enter**. For more information about running NTDSUTIL for Windows Server, see Microsoft's documentation.
 - Type `Activate Instance NTDS`, and then press **Enter**.
 - Type `Authoritative Restore`, and then press **Enter**.
 - Type the following command, and then press **Enter**:
`restore subtree ou=OU_Name,dc=Domain_Name,dc=xxx`
In this command, `<ou_name>` is the name of the organizational unit that you want to restore, `<domain_name>` is the domain name that the OU resides in, and `<xxx>` is the top-level domain name of the domain controller, such as `com`, `org`, or `net`.
 - Repeat these steps as many times as necessary for the specific objects that you need to restore.
 - After you have finished restoring Active Directory information, exit NTDSUTIL.
 - Restart the computer.
7. When you launch Backup Exec, the program prompts you for the database encryption key file.
- Complete the following steps to import the database encryption key file:
- Locate the database encryption key from the secure location to which you backed it up. Backup Exec indicates the name of the key that needs to be restored.
 - Copy the file and then paste it in the Data folder in the directory in which you installed Backup Exec.
8. Log on to Backup Exec.

Forever Incremental Backup

This chapter includes the following topics:

- [About Forever Incremental Backup](#)
- [How do forever incremental backups work](#)
- [Supported storage in forever incremental backups](#)
- [Backing up virtual machines using forever incremental backups](#)
- [CAS-MBES scenarios in forever incremental backups](#)
- [Notes for forever incremental backups](#)
- [Recommendations for forever incremental backups](#)
- [Limitations of forever incremental backups](#)

About Forever Incremental Backup

Customers normally take weekly full backup and daily incremental backups from the source virtual machines. This is commonly known as a traditional virtual-based backup.

Frequent full backups of virtual machines can cause the backups to run past the backup window. This not only affects the SLA, but also affects the virtual machines and the network bandwidth.

Backup Exec introduces Forever Incremental Backup, also known as Backup Exec Accelerator. When you create a forever incremental backup policy, a full backup is taken from the source virtual machine that is followed by incremental backups. A consolidate backup is run by consolidating the previous set of full and

incremental backups. Subsequent incremental backups use the consolidate full backup as a baseline to determine changes in the source virtual machine. A consolidate full backup is equivalent to a full backup from the source virtual machine. It can be used for all types of restores, such as Virtual machine restore, GRT restore, application GRT restore, Instant Recovery, and Recovery Ready.

There are some differences between the traditional and the forever incremental backups of virtual machines.

See [“Differences between the traditional and the forever incremental backups of virtual machines”](#) on page 141.

Benefits of Forever Incremental backups

- Frees up the backup window.
- Meets the backup policy requirements and recovery time objectives by running weekly consolidated backups without running a full backup from the source virtual machine.
- Reduces the load on virtual machines and the network bandwidth by avoiding frequent full backups from the source virtual machine.

Forever incremental backups are supported for both VMware and Hyper-V virtual machines.

Forever incremental also supports multi-stream backup of virtual machines.

See [“Multi-stream backup of virtual machines”](#) on page 1033.

How do forever incremental backups work

Two job templates are added by default for a forever incremental backup definition. You can also add a third job template. You can navigate to the templates from **Backup Options > Schedule** tab.

■ Forever Incremental

Use this template to get data from the source virtual machine with the VMware or Hyper-V backup methods. If it is the first job that runs as a part of the forever incremental solution, it backs up the full data from the source virtual machine.

■ Consolidate Full

Use this template to consolidate all backup sets available since the last full backup and create a new full backup set on the Backup Exec server. The consolidate backup runs locally on the Backup Exec server and does not connect to the virtualization host or the virtual machine that is backed up.

Subsequent incremental backups refer to this consolidate full backup as the baseline.

■ **Full**

Add this job template to run a full backup from the source virtual machine (protect the virtual machine directly from the source) as per the defined schedule.

In a consolidated full backup, at the end of the backup phase, Backup Exec mounts the backup set from the consolidate full backup and compares the GRT volume level information with the last backup set.

For example, you created a forever incremental backup definition for a VMware virtual machine. The incremental is scheduled daily starting from Saturday, 7th September and consolidate full is scheduled weekly starting Friday, 13th September. On the first Saturday, 7th September at 11:00 P.M., full data of the virtual machine is backed up from the source virtual machine (F1).

Table 22-1

	Date	Backup		Date	Backup
Week 1	Saturday, 7th September - 11:00 P.M.	Full (F1)	Week 2	Saturday, 14th September - 11:00 P.M.	Incremental (F2-I2)
	Sunday, 8th September - 11:00 P.M.	Incremental (F1-I1)		Sunday, 15th September - 11:00 P.M.	Incremental (F2-I3)
	Monday, 9th September - 11:00 P.M.	Incremental (F1-I2)		Monday, 16th September - 11:00 P.M.	Incremental (F2-I4)
	Tuesday, 10th September - 11:00 P.M.	Incremental (F1-I3)		Tuesday, 17th September - 11:00 P.M.	Incremental (F2-I5)
	Wednesday, 11th September - 11:00 P.M.	Incremental (F1-I4)		Wednesday, 18th September - 11:00 P.M.	Incremental (F2-I6)
	Thursday, 12th September - 11:00 P.M.	Incremental (F1-I5)		Thursday, 19th September - 11:00 P.M.	Incremental (F2-I7)
	Friday, 13th September - 11:00 A.M.	Full (F2)		Friday, 20th September - 11:00 A.M.	Full (F3)
	Friday, 13th September - 11:00 P.M.	Incremental (F2-I1)		Friday, 20th September - 11:00 P.M.	Incremental (F3-I1)

The backup chain then continues.



Default schedule

You can set a schedule for each job template.

■ Forever Incremental

The default schedule is set for 11:00 P.M. every day. An incremental backup is taken from source virtual machine every day at 11 P.M.

■ **Consolidate Full**

The default schedule is set for every week on Friday at 11:00 A.M. The consolidate operation runs on the Backup Exec server every Friday at 11:00 A.M.

■ **Full**

The default schedule is set for the fourth Friday of each month at 11:00 P.M. A full backup runs on the source virtual machine on the fourth Friday of each month at 11:00 P.M. You can change the schedule as per your requirements.

Data lifecycle management (DLM)

You need to specify the backup set retention for forever incremental backups and consolidate full backup sets. The backup sets are not expired and deleted until the retention period has expired.

Prevent backup sets from getting deleted

Backup sets that are created by forever incremental backup jobs are retained (consolidate full, full, and forever incremental backups). When the consolidate full or full backup from the source virtual machine runs and is successful, the previous backup chain containing the earlier full and the incremental backup sets are no longer retained.

The latest full backup set created by the consolidate full or full backup from the source virtual machine is retained along with the next set of incremental backups. The backup set is retained until the subsequent consolidate full or full backup runs.

Even if a backup set has expired, it is not deleted as it is part of the backup chain that is retained.

After the consolidate full or full backup from the source virtual machine is successfully completed, the sets in the previous chain are no longer retained. Those sets can be deleted depending on the retention period.

Considering the previous example, Backup Exec retains F1 to F1-I5. After F2 runs and is successful, F2 is retained and Backup Exec no longer retains F1 to F1-I5. In this backup chain, F2 to F2-I7 are retained till F3 runs and is successful. At any point, only latest backup chain is retained.

In the Backup Exec user interface, all backup sets have the retention option set as **System defined** and retain reason as **Retained for Consolidate Full job**.

Cataloging and Granular Recovery Technology (GRT)


The consolidate full backup does not run a catalog operation to collect GRT information. The GRT information is instead copied from the latest incremental backup available before the consolidate full backup is run. For the incremental or the full backup from the source virtual machine, the catalog operation is run to collect GRT information.

Supported storage in forever incremental backups

A forever incremental backup definition is supported on the following devices:

- **Disk storage and disk storage that is hosted on a network share**
A consolidate full backup to disk storage reads all the backup sets in the previous backup chain and writes it to a new consolidate backup set. It involves significant read-write operations to create the consolidate backup set. The time taken for this operation and the disk space occupied by the consolidate backup set is proportional to the consolidate backup set size.
- **Deduplication disk storage**
A consolidate full backup to deduplication storage leverages the Backup Exec Deduplication Block Cloning technology. It avoids significant read-write operation to create the consolidate backup set. This new backup set uses references to the blocks present in previous backup sets. The time taken for this operation and the disk space occupied by the consolidate backup set is significantly less compared to a consolidate backup targeted to disk storage.
- **Disk storage on ReFS volumes**
A consolidate full backup to disk storage on ReFS volume leverages Block Cloning technology. The consolidate full job runs faster because it avoids large read-write operation to create the consolidate backup set. The time taken for this operation and the disk space occupied by the consolidate backup set is significantly less compared to a consolidate backup targeted to non-ReFS storage and disk storage that is hosted on a network share. You do not require Backup Exec deduplication disk storage to leverage Block Cloning technology for ReFS volumes.
Prerequisites:
 - The feature only works when the disk storage is configured on an ReFS volume formatted with Windows 2016 or later.
 - Forever incremental backups and consolidate full backup must be on the same disk storage, which supports ReFS Block Cloning.

The following image displays the consolidate backup workflow to Disk storage and how Block Cloning Technology works. The example considers one Forever Incremental backup chain consisting of 1 full backup set, 3 incremental backup sets, and 1 consolidate backup set.



Note: The Duplicate stage can be targeted to all devices that are supported by Backup Exec.

Backing up virtual machines using forever incremental backups

You can create a backup for VMware and Hyper-V virtual machine using the forever incremental backup definition.

To back up virtual machines using forever incremental backups

- 1 On the **Backup and Restore** tab, select a VMware or Hyper-V virtual machine, an ESXi host, or a vCenter server, or a Hyper-V host from the list of servers.
- 2 Do one of the following:

- In the **Backups** group, select **Backup > Create a Forever Incremental Backup for Virtual Machines**.
 - Right-click and select **Backup > Create a Forever Incremental Backup for Virtual Machines**.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit** to add or remove resources from the backup selection list.

Depending on whether you select a VMware ESX host or a Hyper-V host, the virtual machines for the selected host are displayed.
 - 4 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**, and select the backup options.
 - 5 On the **Backup Definition Properties** dialog box, click **OK**.

CAS-MBES scenarios in forever incremental backups

Review the following information that is related to CAS-MBES.

- In a CASO environment, it is recommended that forever incremental, consolidate full, and full jobs are targeted to storage devices from the same host. This helps to avoid an impact of network bandwidth on the performance of the consolidate full backup job.
- In a CASO environment, for **Backup to disk**, you can select multiple disk storage across various servers.
- In a CASO environment, there is a restriction for the deduplication device to ensure that no performance impact is observed. Forever incremental backups allow selection of only one deduplication device for the jobs in the **Backup** stage. You can only select the host server of the deduplication disk storage as a Backup Exec server.
- In a rolling upgrade scenario, forever incremental backup definition jobs can only run on Backup Exec 20.6 or later.

Notes for forever incremental backups

The following notes are for forever incremental backups.

- To leverage the Block Cloning technology for faster consolidated backups, target forever incremental and consolidated backup jobs to Backup Exec deduplication disk storage or to a disk storage on ReFS volume only. Duplicate jobs can be targeted to any storage.

- If the configuration for a VMware or Hyper-V virtual machine changes, the next incremental backup is promoted to a full backup. For example, a new disk is added to a virtual machine and if the next backup is an incremental backup, it gets promoted to a full backup.
- To create a forever incremental backup definition, create a new job. Traditional backup definitions cannot be converted into forever incremental backup definition.
- The **Run a SQL Log backup after backing up the virtual machine** setting for SQL Application GRT virtual machine backup is not applicable to the consolidate full job. The setting is applicable to the full or the forever incremental backup from the source virtual machine.
- For Hyper-V virtual machines to leverage Block Cloning Technology on ReFS volumes, ensure that the sector size of the virtual disk matches the cluster size of the ReFS volume on the media server on which the backup is being taken.

Due the limitations of ReFS Block Cloning technology, the following behavior may be observed:

- Unaligned source and destination blocks require significant read-write operations.
- Block Cloning technology is not leveraged optimally.
- Performance impact on the consolidate full backup.

For more information about the limitations of using block cloning on ReFS, refer to the Microsoft documentation.

Recommendations for forever incremental backups

Recommendations to help you use forever incremental backups.

- Improve consolidate full backup performance significantly by targeting the job to Backup Exec deduplication disk storage or to a disk storage on ReFS volume (Windows 2016 or later) and leverage Block Cloning technology.
- Consolidate backup is an I/O intensive operation if Block Cloning technology is not used. Ensure that you select the proper backup storage for Forever Incremental backup definition.
- It is recommended that a Consolidate Full backup runs after a maximum of 30 consecutive incremental backups. This ensures that restore operations complete in a reasonable time.
- It is recommended that you run a periodic full backup from the source virtual machine. Compared to traditional backups, in a forever incremental backup

definition, a full backup from the source virtual machine can be run at a lesser frequency. This is applicable when a virtual machine has significant configuration changes.

For example, full backup from the source virtual machine - monthly, incremental - daily, and consolidate full - weekly.

- It is recommended that you schedule a consolidate full backup, outside the regular backup window.
- Keep the **Verify** option enabled for forever incremental backup definition.
- If a consolidate full job fails, it is recommended that you run a full backup from the source virtual machine.
- Run the **Validate VM** job after a consolidate full set. You can run this job when backup sets are hosted on a disk.

Limitations of forever incremental backups

Limitations when you use forever incremental backups.

- Forever incremental backups for Hyper-V are supported only using the Standard processing and Resilient Change Tracking (RCT) method. You cannot run a forever incremental backup using the Faster processing method.
- Disk storage on ReFS volumes for VMware and Hyper-V machines:
 The volume must be formatted with ReFS on Windows Server 2016 or later. If Failover Clustering is in use, at the time of formatting, the Clustering Functional Level must be at Windows Server 2016 or later.
 For more information about the limitations of using block cloning on ReFS volumes, refer to the Microsoft documentation.

Backup Exec Agent for Windows

This appendix includes the following topics:

- [About the Agent for Windows](#)
- [Requirements for the Agent for Windows](#)
- [Stopping and starting the Agent for Windows](#)
- [Establishing a trust between the Backup Exec server and a remote computer](#)
- [About the Backup Exec Agent Utility for Windows](#)
- [Using the Backup Exec Agent Utility Command Line Applet](#)
- [Backup Exec Agent Utility Command Line Applet switches](#)

About the Agent for Windows

The Agent for Windows enables Windows Servers network administrators to perform backup and restore operations on Backup Exec agents that are connected to the network.

The Agent for Windows is a system service that runs on remote Windows servers and workstations. The Agent for Windows provides faster backup processing by locally performing the tasks that require extensive network interaction in typical backup technologies. The Agent for Windows processes backup data into a continuous stream that the Backup Exec server then processes as a single task. This method provides better data transfer rates over traditional technologies, which require multiple requests and acknowledgments between the Backup Exec server and the remote server.

The Agent for Windows enables you to do the following:

- Back up and restore in firewall environments.
- Back up and restore using a specified local network if the Backup Exec server and the remote computer are on the same subnet.
- Attain significant performance increases when running modified backups (for example, differential and incremental). This occurs because file selection is performed locally by the Agent for Windows instead of across the network as performed by traditional network backup applications.

Note: Network hardware has a major effect on performance. Performance is directly related to the capabilities of the networking hardware in the Backup Exec server and the remote device. Higher network bandwidth ratings also contribute to faster operation processing.

See [“Requirements for the Agent for Windows”](#) on page 952.

See [“Methods for installing the Agent for Windows”](#) on page 68.

See [“Changing network and security options for Backup Exec”](#) on page 718.

See [“Using Backup Exec with firewalls”](#) on page 725.

See [“Backup Exec Shadow Copy Components file system”](#) on page 188.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

Requirements for the Agent for Windows

Because an Agent for Windows is also a Client Access License (CAL), you must install the Agent for Windows on any remote Windows computer that you want to back up. You cannot fully protect resources on a remote server until an Agent for Windows has been installed.

At the Backup Exec server, you must enter Agent for Windows licenses for each remote Windows computer that you want to protect. To back up a remote Windows computer from more than one Backup Exec server, you must enter the same Agent for Windows licenses on each Backup Exec server.

Backup Exec Agent for Applications and Databases also includes an Agent for Windows that lets you protect one remote Windows computer. The Agent for Windows license is enabled when you install the database agents on the Backup Exec server.

To protect the Workstation versions of the supported Windows platforms, you must install the Agent for Windows on each platform.

You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

Note: If a previous version of the Agent for Windows is installed, it is automatically upgraded when you initiate a new Agent for Windows installation. Previous versions of the Agent for Windows are automatically detected on the remote computers and replaced with the new version during installation of the new Agent for Windows. The name of the system service may have changed when the upgrade is complete.

You can install the Agent for Windows using many methods, depending on your environment.

See [“Methods for installing the Agent for Windows”](#) on page 68.

Stopping and starting the Agent for Windows

The Agent for Windows is automatically started as a service when Windows is started on the remote computer.

To stop or start the Agent for Windows

- 1 Open Windows Services.
- 2 In the **Results** pane, right-click **Backup Exec Remote Agent for Windows**.
- 3 Do one of the following:

To stop the Agent for Windows	Click Stop .
To start the Agent for Windows	Click Start .

See [“About the Agent for Windows”](#) on page 951.

Establishing a trust between the Backup Exec server and a remote computer

When you connect to a remote computer from the Backup Exec server, you must establish a trust between the Backup Exec server and the remote computer to ensure secure communication. You must also establish this trust if you want to configure a remote computer to perform client-side deduplication. You should manually verify the identity of the remote computer to ensure that the remote computer is a trusted source before you establish the trust. After you verify that

the remote computer is a trusted resource, you can establish the trust with the Backup Exec server.

Backup Exec issues a security certificate for both the Backup Exec server and the remote computer. The security certificate is valid for approximately one year and is automatically renewed during normal operations. However, if the certificate expires you must re-establish the trust.

You can establish a trust between the Backup Exec server and the remote computer by adding the remote computer to the list of servers on the **Backup and Restore** tab.

To establish a trust for a remote computer

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Click Microsoft Windows computers and servers.
- 3 Follow the on-screen prompts.

About the Backup Exec Agent Utility for Windows

The Backup Exec Agent Utility is installed when the Agent for Windows is installed on a remote Windows computer.

You can perform the following tasks with the Backup Exec Agent Utility:

- Start the Backup Exec Agent Utility each time you log on.
See [“Starting the Backup Exec Agent Utility”](#) on page 955.
- View current activity on the remote Windows computer.
See [“Viewing the activity status of the remote computer from the system tray”](#) on page 956.
- Configure the Agent for Windows to send information such as the version and the IP address to a Backup Exec server.
See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.
- Configure the Backup Exec Agent Utility for backup and restore operations of Oracle instances.
See [“Configuring an Oracle instance on Linux servers”](#) on page 1235.
- Configure the Backup Exec Agent Utility for Backup Exec server database access for Oracle operations.
See [“Configuring database access for Oracle operations”](#) on page 960.
- Remove the security certificate for a Backup Exec server.

See [“Removing Backup Exec servers that the Agent for Windows can publish to”](#) on page 960.

Starting the Backup Exec Agent Utility

You access the Backup Exec Agent Utility from the Windows taskbar.

See [“Viewing the activity status of the remote computer in the Backup Exec Agent Utility”](#) on page 955.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.

To start the Backup Exec Agent Utility

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click the icon to view the utility.

- 2 To open the registry editor, the Services window, and the Event Viewer on the remote Windows computer, right-click the Backup Exec Agent Utility icon in the system tray, and then click **Tools**.

Viewing the activity status of the remote computer in the Backup Exec Agent Utility

You can use the Backup Exec Agent Utility to view the activity status of the remote Windows computer.

To view the activity status of the remote computer in the Backup Exec Agent Utility

- 1
- On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2
- Click the **Status** tab.

You can view any of the following information about the remote Windows computer:

Backup Exec server	Displays the name of the Backup Exec server that is processing the current operation.
Source	Displays the media or share that is being processed.
Current folder	Displays the name of the current directory, folder, or database (depending on the specific agent) that is being processed.
Current file	Displays the name of the current file that is being processed.

- 3
- Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

Viewing the activity status of the remote computer from the system tray

You can view the activity status for a remote computer.

Possible statuses are as follows:

- A backup job is running.
- A restore job is running.
- A backup and a restore job are running.
- Snapshot in progress.
- The Backup Exec client service named Beremote.exe is not running on the computer.
- Idle.

To view the activity status of a remote computer

- ◆ Position the cursor over the Agent for Windows icon in the system tray.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

Starting the Backup Exec Agent Utility automatically on the remote computer

You can start the Backup Exec Agent Utility automatically each time you log on to the remote computer.

To start the Backup Exec Agent Utility automatically on the remote computer

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.
- 3 Check the **Start the Backup Exec Agent Utility every time you log on** check box.
- 4 Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

Setting the refresh interval on the remote computer

You can display the number of seconds for the Backup Exec Agent Utility to wait before refreshing the status of the computer.

To set the refresh interval on the remote computer

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.
- 3 In the **Refresh interval** box, type the number of seconds to refresh the status.
- 4 Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

About publishing the Agent for Windows to Backup Exec servers

Use the Backup Exec Agent Utility to add, change, or delete the Backup Exec server names or IP addresses that this remote Windows computer publishes to. Each Backup Exec server that you add to the list on the **Publishing** tab displays the remote computer in the list of servers on the **Backup and Restore** tab. After the remote computer is added to the list of servers, you can right-click the remote computer and click **Establish Trust** to verify that the remote computer is a trusted resource.

This information that the Agent for Windows publishes includes the version of the Agent for Windows and the remote computer's IP addresses. Because the remote computer's IP address is published to the Backup Exec server, the Backup Exec server can connect to and display the remote computer even if it is in an unknown domain.

For each Backup Exec server that is published to, you can specify a local backup network for operations between the Backup Exec server and the remote computer. Directing jobs to a specified local network rather than to a corporate network isolates the backup data traffic so that other connected networks are not affected when operations are performed between the Backup Exec server and the remote computer.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

See [“Backup networks”](#) on page 716.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

See [“Adding Backup Exec servers that the Agent for Windows can publish to”](#) on page 958.

See [“Editing Backup Exec server information that the Agent for Windows publishes to”](#) on page 959.

See [“Removing Backup Exec servers that the Agent for Windows can publish to”](#) on page 960.

Adding Backup Exec servers that the Agent for Windows can publish to

You can use the Backup Exec Agent Utility to add a Backup Exec server that the Agent for Windows can publish information.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.

See [“Viewing the activity status of the remote computer from the system tray”](#) on page 956.

To add Backup Exec servers that the Agent for Windows can publish to

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Click **Add**.
- 5 Enter the following information:

Backup Exec server name or IP address	Type the Backup Exec server name or the IP address of the Backup Exec server to which you want to publish information.
User Name	Type the user name for an account that has administrative rights on the Backup Exec server.
Password	Type the password for an account that has administrative rights on the Backup Exec server.

- 6 Click **OK**.

Editing Backup Exec server information that the Agent for Windows publishes to

You can use the Backup Exec Agent Utility to edit a Backup Exec server name or IP address to which the Agent for Windows can publish information.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.

To edit Backup Exec server information

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.

- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to edit from the list.
- 5 Click **Edit**.
- 6 Edit the Backup Exec server name or IP address.
- 7 Click **OK**.

Removing Backup Exec servers that the Agent for Windows can publish to

You can use the Backup Exec Agent Utility to remove a Backup Exec server so that the Agent for Windows no longer publishes information to it.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.

To remove Backup Exec servers that the Agent for Windows can publish to

- 1 the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to remove from the list.
- 5 Click **Remove**.
- 6 Click **OK**.

Configuring database access for Oracle operations

You can configure database access to enable the Backup Exec server to authenticate Oracle operations.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1241.

To configure database access for Oracle operations

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.
- 2 Click the **Database Access** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.

4 Complete the appropriate options to configure database access:

Enable the Backup Exec server to authenticate Oracle operations	Select this option to enable Oracle operations between the Backup Exec server and this computer.
User name	Specify a user name that has administrative rights to this computer. This logon account is what the Backup Exec server uses when it connects to this computer. If you specify an IP address or a fully qualified computer name as part of the user name, the Backup Exec Agent Utility may not be able to verify the user account. If the credentials entered are incorrect, the error “cannot attach to a resource” may be displayed when you run a backup or restore job. You must add this computer name and logon account to the Backup Exec server's list of authentication credentials for Oracle servers. If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you are browsing the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.
Password	Specify the password for this logon account. Note: For security reasons, the logon credentials are not stored on the remote computer.
Confirm Password	Type the password again to confirm it.
Use a custom port to connect to the Backup Exec server during Oracle operations	Select this option to change the port that is used for communications between this computer and the Backup Exec server during Oracle operations. By default, port 5633 is used. If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server.
Port number	Type the port number that you want to use for communications between this computer and the Backup Exec server.

- 5 Click **OK**.
 - 6 On the Backup Exec server, add the name of the Oracle server and the user name that you entered on the **Database Access** tab to the Backup Exec server's list of authentication credentials.
- See [“About the Backup Exec Agent Utility for Windows”](#) on page 954.

Removing a security certificate for a Backup Exec server that has a trust with the Agent for Windows

You can remove the security certificate for a Backup Exec server that has established a trust with the Agent for Windows.

To remove a security certificate for a Backup Exec server

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.
- 2 Click the **Security** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to remove the security certificate for, and then click **Remove**.
- 5 Click **OK**.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 953.

Using the Backup Exec Agent Utility Command Line Applet

You can use the Backup Exec Agent Utility Command Line Applet from any Windows operating system command prompt to access the Backup Exec Agent Utility. The Backup Exec Agent Utility Command Line Applet is installed when you install the Agent for Windows.

If you run the command line utility on a Windows computer, you must run it in elevated command prompt.

Note: To run the Backup Exec Agent Utility Command Line Applet on a Microsoft Windows computer, you must use Server Core.

You can run the following Backup Exec Agent Utility functions with the Backup Exec Agent Utility Command Line Applet:

- Set the publishing interval (in minutes).
- List the published name for the agent.
- List the Backup Exec server names to which the agent is publishing.
- Add a Backup Exec server to the publishing list.
- Remove a Backup Exec server from the publishing list.
- View the following status information:
 - Activity status
 - Current source
 - Current folder
 - Current file
 - Currently attached Backup Exec server

To use the Backup Exec Agent Utility Command Line Applet

- 1 Open a command prompt.
- 2 From the Backup Exec installation directory, type `ramcmd.exe` followed by a series of command switches.

The default installation location is `c:<Backup Exec install path>\Backup Exec\RAWS`

See [“Backup Exec Agent Utility Command Line Applet switches”](#) on page 964.

Backup Exec Agent Utility Command Line Applet switches

The following table describes the switches that you can use with the Backup Exec Agent Utility Command Line Applet.

See [“Using the Backup Exec Agent Utility Command Line Applet”](#) on page 963.

Table A-1 Backup Exec Agent Utility Command Line Applet switches

Switch	Description
status:[n]	Status output is repeated every <n> seconds, with a range of 1 - 86400. Press Q to stop the output from running. <pre>ramcmd /status:[n]</pre> When you use the /status switch without a time value, the Agent for Windows status appears in the command window and then the applet exits.
/publish:[on off add remove interval][/ms:<Backup Exec server>] [/t:<x>]	Use the following parameters with the /publish switch: ■ No parameter specified- Displays the publishing status and then exits. ■ [on] - Turns publishing on. Lets the Agent for Windows send information about itself, such as its version and IP address. ■ [off] - Turns publishing off. ■ [add], [remove] - Used with /ms. You can use this parameter to add or remove Backup Exec servers from the Agent for Windows publish list. ■ [interval] - Used with /t. Specifies the time interval that the Agent for Windows sends information about itself to the Backup Exec server. You can set the time interval in minutes using the /t:[<x>] parameter. Note: The [interval] switch must be used with the /t: switch. Using [interval] alone on the command line is not supported. <pre>ramcmd /publish:[on off add remove interval] [/ms<Backup Exec server>] [/t:<x>]</pre>

Table A-1 Backup Exec Agent Utility Command Line Applet switches
(continued)

Switch	Description
/oracle: [new edit delete] /in:[<instance name>] /ms:[<Backup Exec server address>] /jt:[<job template>] /user:[<username>] /password:[<password> *] /rc: [yes no] /tns:[<TNS name>]	Use the following parameters with the /oracle switch: ■ No parameter specified- Displays the existing Oracle instances and then exits. ■ [new], [edit], [delete] - Used with switch /in. ■ /in:[<instance name>] - Used to add, edit, and delete Oracle instance names from the Oracle instance list. ■ /ms:[<Backup Exec server name address>] - Sets the Backup Exec server name or its IP address. ■ /jt:[<job template>] - Sets a Backup Exec job template. ■ /user:[<username>] - Sets a username. ■ /password:[<password> *] - Sets a password to be used with /user:[<username>]. If you omit the password, or you use an asterisk [*], you do not need to enter the password on the command line. After the command runs, a prompt appears asking you for a password. ■ /rc:[yes no] - Turns the Use recover catalog setting on or off. If /rc appears without a parameter, then the current status for that instance is displayed. ■ /tns:[TNS name] - Sets the TNS name alias of an available Oracle database and the server it resides on in the Oracle TNSNAMES file. <pre> ramcmd.exe /oracle:edit /in:<instance name> /rc:[yes no] [/tns:<TNS name>] [/user:<username>] [/password:password *] </pre>
/auth:[on off] [/user:<username>] [/password:<password> *]	Enables or disables Backup Exec server authentication for Oracle operations. ■ /auth:on - Turns the state on. Requires /user parameter. ■ /auth:off - Turns the state off. Requires /user parameter. ■ /user:<username> - Sets a username. ■ /password:<password> - Sets a password to be used with /user:<username>. If you enter an asterisk for the password or omit the password, you are prompted for the password.

Table A-1 Backup Exec Agent Utility Command Line Applet switches
(continued)

Switch	Description
/port:[<port>]	Displays or sets a custom port that is used to connect to the Backup Exec server during Oracle operations. ■ /port - Displays the current port number. If the port is the default port, displays "(default)". ■ /port:<port> - Sets the port number to <port>. To change the port to the default port number, type [/port:0].
/log_path:[<log path>]	Displays or sets a custom path for debug logs. ■ /log_path - Displays the log directory path and then exits. ■ /log_path:<"logs path"> - Creates the directory <"logs path">. If the path has a space in the name, enclose the path in quotes. For example, "C:\Program files\LogsFolder".

Backup Exec Deduplication Feature

This appendix includes the following topics:

- [About the Deduplication feature](#)
- [Deduplication methods for Backup Exec agents](#)
- [Requirements for the Deduplication feature](#)
- [Configuring a 125 TB Deduplication storage folder in Backup Exec](#)
- [Installing the Deduplication feature](#)
- [Converting an older version of Deduplication Storage to a newer version](#)
- [Creating or importing deduplication disk storage](#)
- [Selecting storage devices for direct access sharing](#)
- [Changing the location of a deduplication disk storage device](#)
- [Sharing a deduplication device between multiple Backup Exec servers](#)
- [How to use client-side deduplication](#)
- [How to set up backup jobs for deduplication](#)
- [Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices](#)
- [Copying deduplicated data to tapes](#)
- [Using deduplication with encryption](#)
- [Restoring a deduplication disk storage device or deduplicated data](#)

- [Disaster recovery of deduplication disk storage devices](#)
- [Disaster recovery of OpenStorage devices](#)

About the Deduplication feature

The Backup Exec Deduplication feature supports a data-reduction strategy by optimizing storage and network bandwidth. The Deduplication feature supports integrated deduplication at the Backup Exec server and on remote computers that have the Agent for Windows or the Agent for Linux installed. It also allows data to be deduplicated and stored on intelligent disk devices from other vendors.

Table B-1 Types of deduplication

Type of deduplication	Where deduplication occurs	Benefits
Backup Exec server-side deduplication	On the Backup Exec server.	Reduces the size of backups, which reduces storage requirements.
Client-side deduplication	On the remote computer where the data is located. Note: The Agent for Windows is required on the remote Windows computer to perform Windows client-side deduplication. The Agent for Linux is required on the Linux computer to perform Linux client-side deduplication.	Reduces network traffic because only unique data is sent across the network. It also reduces the backup window.
Appliance deduplication	On an intelligent disk device, such as Arctera PureDisk, or a device from a third-party vendor.	Reduces the size of backups, which reduces storage requirements. It also reduces the backup window.

With a single Deduplication feature license key, you can use two types of deduplication devices.

Table B-2 Types of deduplication devices that work with the Deduplication feature

Type of device	Description
OpenStorage device	Backup Exec uses Veritas's OpenStorage technology, which allows intelligent disk devices to integrate with Backup Exec. You can back up data to the Arctera PureDisk device and to storage devices from other vendors. You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List. You can find more information about OpenStorage devices, prerequisites, configuration, and properties. See “OpenStorage devices” on page 418. See “Configuring an OpenStorage device” on page 419.
Deduplication disk storage	Deduplication disk storage provides integrated deduplication on the Backup Exec server. Deduplication disk storage is a disk-based backup folder that is located on the Backup Exec server. See “Creating or importing deduplication disk storage” on page 982.

In addition to reducing storage requirements and network traffic, the Deduplication feature lets you do the following:

- Copy deduplicated data from an OpenStorage device or deduplication disk storage to tape for long-term or off-site storage.
- Use optimized duplication, which lets you copy deduplicated data between OpenStorage devices from the same vendor and between deduplication disk storage devices.
- Use Granular Recovery Technology (GRT) with jobs that use deduplication devices.
- Share OpenStorage devices and deduplication storage devices among multiple Backup Exec servers when you use the Central Admin Server feature.

For information about the best practices to use Backup Exec Deduplication and Backup Exec Deduplication with the Central Admin Server Option (CASO), refer to *Backup Exec Best Practices*.

See [“Installing the Deduplication feature”](#) on page 980.

See [“Requirements for the Deduplication feature”](#) on page 972.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 995.

See [“Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices”](#) on page 997.

See [“Copying deduplicated data to tapes”](#) on page 1000.

Deduplication methods for Backup Exec agents

Backup Exec supports the following deduplication methods:

- Client-side deduplication, either on an intelligent disk device or to a deduplication disk storage device.
- Backup Exec server-side deduplication with a deduplication disk storage device.
- Appliance deduplication on an OpenStorage device.

The following table lists the deduplication methods that are available for the Backup Exec agents.

Table B-3 Deduplication methods for Backup Exec agents

Agent	Client-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Client-side deduplication (with Granular Recovery Technology enabled)	Backup Exec server-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Backup Exec server-side deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Agent for Windows	Yes	Not applicable	Yes	Not applicable	Yes
Agent for VMware and Hyper-V	Yes (for Hyper-V only) Note: The Agent for Windows must be installed on the Hyper-V host.	Yes (for Hyper-V only) Note: The Agent for Windows must be installed on the Hyper-V host.	Yes	Yes	Yes

Table B-3 Deduplication methods for Backup Exec agents (*continued*)

Agent	Client-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Client-side deduplication (with Granular Recovery Technology enabled)	Backup Exec server-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Backup Exec server-side deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Agent for Linux	Yes	No	Yes	Not applicable	Yes
Agent for Enterprise Vault	No	No	Yes	No	No
Exchange Agent	Yes	Yes	Yes	Yes	Yes
SQL Agent	Yes	Not applicable	Yes	Not applicable	Yes
SharePoint Agent	Yes	Yes	Yes	Yes	Yes
Active Directory Agent	Yes	Yes	Yes	Yes	Yes
Agent for Oracle	Linux: Yes Windows: Yes	No	Yes	No	Yes

See [“About the Deduplication feature”](#) on page 969.

Requirements for the Deduplication feature

The requirements for the Deduplication feature vary depending on the type of storage devices you want to use and the type of deduplication you want to use. Before you use the Deduplication feature, you should determine what type of storage devices you want to use with it and what type of deduplication you want

to use. Then, verify that your system meets the requirements for the storage devices you want to use.

Warning: It is strongly recommended that you exclude the deduplication disk storage device from all antivirus scans. If an antivirus scanner deletes or quarantines the files from the deduplication disk storage device, access to the device may be disabled.

Table B-4 Requirements and recommendations for the Deduplication feature

Item	Requirements and Recommendations
Deduplication disk storage devices	

Table B-4 Requirements and recommendations for the Deduplication feature
(continued)

Item	Requirements and Recommendations
	The following items are required: ■ A 64-bit Backup Exec server. ■ A Backup Exec server with a minimum of 8 cores. It is recommended that you have 16 cores. ■ For 100 TB of stored deduplicated data, 16 cores are required, 32 cores are recommended. ■ A dedicated volume to use as the location to store the deduplication disk storage. The dedicated volume must have 20 percent free space that totals no less than 5 gigabytes (GB) of free space. ■ The deduplication disk storage device must be excluded from antivirus scans. If an antivirus scanner deletes or quarantines the files from the deduplication disk storage device, access to the deduplication disk storage device may be disabled. ■ Deduplication disk storage requires 8 GB physical memory for up to 4 TB storage. After that 1 GB of physical memory is required for each 1 TB of deduplication disk storage and up to 32 TB. For example, 9 GB physical memory for 5 TB storage. For deduplication disk storage more than 32 TB and up to 64 TB, 32 GB or more physical memory is recommended. For deduplication disk storage more than 64 TB and up to 100 TB, 128 GB or more physical memory is recommended. ■ It is recommended the following as minimum disk speeds per individual read, write, or verify operation: ■ Up to 32 TBs of storage: ■ 130 MB per second ■ 200 MB per second for enterprise-level performance ■ 32 to 48 TBs of storage: 200 MB per second ■ 48 to 64 TBs of storage: 250 MB per second ■ 64 to 100 TBs of storage: 350 MB per second The above recommendations are for the performance of a single operation. You may need more capability depending on your objectives for writing to and reading from disk. Note: You should be aware of the effects that computer disk speeds have on deduplication performance. Computer disk speeds can have the following effects on deduplication performance: ■ Computers with disk speeds greater than 200 MB per second

Table B-4 Requirements and recommendations for the Deduplication feature
(continued)

Item	Requirements and Recommendations
	have optimal read and write performance for deduplication. ■ Computers with disk speeds between 150-200 MB per second have sufficient read and write speed for deduplication. ■ Computers with disk speeds between 100-150 MB per second have some operations with degraded performance. ■ Computers with disk speeds less than 100 MB per second experience poor performance. You should improve disk reads and writes before you install and run deduplication. See “Configuring a 125 TB Deduplication storage folder in Backup Exec” on page 977.
Deduplication disk storage connection	■ Storage area network (Fibre Channel or iSCSI), direct-attached storage (DAS), or internal disks are supported. Removable disks including USB, eSATA, and FireWire devices are not supported. ■ The Backup Exec server should have redundant connectivity to the storage. ■ The storage network must be a dedicated, low latency network with a maximum of 1-millisecond latency per round trip. ■ The storage network must have enough bandwidth to meet your throughput objectives. The following storage network bandwidths are supported: ■ iSCSI SANs with a bandwidth of at least 10 Gb per second. ■ Fibre Channel SANs with a bandwidth of at least 4 Gb per second. ■ A minimum bandwidth of 130 MB per second is required for read and write performance. Bandwidth that is less than 130 MB per second may be used in smaller, less resource-intensive environments. However, as usage increases, deduplication requires more bandwidth to ensure adequate throughput for deduplication processes and backups. Otherwise, performance and stability are negatively affected.
Deduplication disk storage credentials	The following requirements are for the password credentials for a deduplication disk storage device: ■ The password cannot be blank ■ The password cannot contain the following characters: &, ", <, >, %, ^ ■ The password cannot start with the hyphen character: - ■ The password cannot end with the backslash character: \

Table B-4

Requirements and recommendations for the Deduplication feature
(continued)

Item	Requirements and Recommendations
OpenStorage devices	To use a Arctera PureDisk device or a storage device from another vendor as an OpenStorage device, you must purchase the device and the appropriate OpenStorage connector from the device's vendor. You can use the Deduplication feature with OpenStorage devices on a 64-bit Backup Exec server. The standard system requirements for Backup Exec apply to the Deduplication feature when you use OpenStorage devices. You can find more information about OpenStorage devices, prerequisites, configuration, and properties. See “OpenStorage devices” on page 418.
Client-side deduplication for Windows	On the server where the Agent for Windows is installed, 1.5 GB of memory is required. Both 32-bit and 64-bit Windows operating systems are supported.
Client-side deduplication for Linux	You can find a list of compatible operating systems for Linux client-side deduplication in the Backup Exec Software Compatibility List. A 64-bit Linux operating system is required. The following deduplication devices can be used: ■ Deduplication disk storage device ■ Arctera PureDisk OpenStorage device, which is the only type of OpenStorage device that supports client-side deduplication for Linux. The following Backup Exec options are required: ■ Agent for Linux ■ Deduplication feature

See [“Installing the Deduplication feature”](#) on page 980.

Configuring a 125 TB Deduplication storage folder in Backup Exec

Backup Exec supports expansion of the Deduplication folder up to 125 TB. This section helps you to configure the 125 TB Deduplication storage folder in Backup Exec.

The maximum supported capacity of the Deduplication folder is 125TB out of which 20% is required for deduplication database and internal log processing. So effectively maximum 100 TB of backup data can be stored.

Minimum hardware requirements

- **CPU:** A 64-bit processor with a minimum clock rate of 2.4-GHz is required. A minimum of 8 cores are required, 16 cores are recommended.
- **Memory:** Minimum 125 GB. You may need to add more memory if you have additional roles performed by the same media server.
- **Storage:**
 - **Metadata disk:** RAID 0+1 is recommended, with at least 1TB of space.
 - **Data disks:** It is recommended to have a maximum of five mount points. Each mount point must have a separate RAID group, a RAID 6 is recommended. Both the metadata disk and data disk must have more than 250 MB/sec of read or write speed.
 - Disks including more number of spindles provide better performance. Storage controller cache of minimum 2 GB or more is recommended.

Notes for configuring the 125 TB Deduplication storage folder

- The Ransomware Resilience feature currently does not protect additional data partitions created using this configuration.
- If you do not follow the settings, there can be performance issues as data may not be balanced optimally across all volumes.
- Support for Deduplication folder size is up to 125 TB. The Deduplication folder can be of a smaller size and you can later add a maximum of five additional data partitions.

To configure 125 TB deduplication storage folder on a Windows server using 25 TB volumes

- 1 Configure a Deduplication storage folder on local storage on Backup Exec server.
- 2 Disable lockdown setting from the Backup Exec global settings.
Perform the steps in the order listed:
 - Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
 - In the left pane, select **Network and Security**.
The **Disk storage lockdown setting**, is enabled by default.
 - Click **Disable**, to disable the setting.

At the prompt, enter the System Logon Account password and the reason for disabling the setting.

Note: Lockdown setting ensures that access to Backup Exec disk storage is limited only to authorized processes, such as Backup Exec services. Only Backup Exec is allowed to write to the disk storage. Lockdown is disabled to ensure that commands listed in the following step are allowed to temporarily write to the storage. To prevent data loss, this setting must be enabled as soon as the configuration is completed.

- 3 Create data folder volumes under Deduplication storage as data1, data2, data3, data 4, data5 of 25 TB each.

Backup Exec qualified this solution only using NTFS volumes. Volumes must be mounted to an empty folder using disk management before running the following commands.

For more information about mounting a volume in an empty folder, refer to the following link:

<https://docs.microsoft.com/en-us/windows-server/storage/diskmanagement/assignmountpointtoinactive>

Perform the following in the order listed:

- Create an `etc` directory and `nbapp-release` file using the following commands:

```
mkdir c:\etc
echo Windows_BYO > "c:\etc\nbapp-release"
```

- Update `DCHeaderHashSize` setting in `<Dedupe Storage DIR>\etc\puredisk\contentrouter.cfg` to **2000000 / number_of_volumes**.

For example, if you have five mount points, set `DCHeaderHashSize` to **400000**.

- Create Additional Data Partitions for Deduplication using the following command:

```
C:\Program Files\Veritas\Backup Exec>crcontrol.exe
--dsaddpartition E:\BackupExecDeduplicationStorageFolder\data1
```

Run the earlier command for the remaining data partitions.

- 4 After all the commands complete successfully, restart the Backup Exec services including the Deduplication services.

Additional Data mount points are available inside the Deduplication folder.

- 5 Enable lockdown setting from the Backup Exec global settings.

Installing the Deduplication feature

Before you attempt to install a Backup Exec edition which includes the Deduplication feature, verify that your system meets the requirements.

See [“Requirements for the Deduplication feature”](#) on page 972.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

Converting an older version of Deduplication Storage to a newer version

Backup Exec improves speed and deduplication when backing up to Backup Exec deduplication folder. The conversion from an older version of deduplication storage to a newer version involves the conversion of the existing deduplication data to a newer format. The conversion time depends on the size of the deduplication storage and the number of backup sets.

To convert from an older version of Deduplication Storage to a newer version

- 1 It is recommended that you make a copy of the deduplication data before the upgrade starts.
- 2 Ensure that the following are true:
 - You have at least 12% free space available on the volume on which the deduplication storage folder exists.
 - Deduplication services are running.
 - Windows Hotfix is installed.

See [“Running the Environment Check before installing or upgrading Backup Exec”](#) on page 42.

- 3 If you need to free up more space on the volume, you can use the tool `pdde_gc.exe` to reclaim storage space. If you have enough free space available, proceed to the next step.

To run `pdde_gc.exe`, follow these steps:

- Mount the ISO from the Backup Exec media or upgrade an earlier version of Backup Exec. On the media, the tool is located in
`<mounted-path>\BE\WinNT\Install\PDDEMigration.`
- On the command line, run `pdde_gc.exe` without any parameters.

If you need more free space, you can try any of the following options:

- Use the Backup Exec console to expire backup sets on the deduplication storage.
 - Run CR queue processing twice.
 - Run the pdde_gc.exe tool again.
- 4 Upgrade Backup Exec to the latest version.

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 48.

After Backup Exec has been upgraded, the conversion from the older version of deduplication storage to a newer version starts. The existing deduplication data is converted to a new format. During the conversion process, the deduplication storage remains offline. Any ongoing deduplication jobs fail and jobs targeted to any other storage continue to run during the deduplication storage conversion.

The Backup Exec console displays an alert that the conversion process has started. Depending on the time required for the conversion, an alert is displayed every 15 minutes showing the progress of the conversion. An alert is also displayed if the conversion is successful, has resumed, or has failed.

- 5 If the conversion is successful, a prompt is displayed in the Backup Exec console, asking you to restart the Backup Exec services.

If you click **OK**, the Backup Exec Services dialog box is displayed. Click **Restart all services**.

Optionally, on the Backup Exec Administration Console, click the Backup Exec button, select **Configuration and Settings**, select **Backup Exec Services**, and then click **Restart all services**.

If the conversion fails, you can manually convert the data to the new deduplication format. Refer to the following section for information on how to manually convert the data.

If the conversion fails because of a server restart, the conversion resumes after the service restart is complete.

Create or import an older deduplication folder

You can create or import a deduplication storage folder in Backup Exec. You have an existing deduplication storage folder that was created when you had the older deduplication version that Backup Exec no longer supports. When you try to import the folder, an error is displayed and the import fails. You need to manually convert the older version of the folder to the new supported deduplication version supported by Backup Exec using the steps in the following section:

After the conversion is completed successfully, you can import the folder. There can be only one deduplication storage configured with a single Backup Exec media server.

Restore an older deduplication storage backup using Backup Exec

You can restore an existing deduplication backup set using Backup Exec. This backup set was created and backed up using an older deduplication version that Backup Exec no longer supports. When you try to restore the backup set, the restore job fails with an exception regarding the older deduplication backup sets but the data is restored. You need to manually convert the restored deduplication folder to a newer supported version of deduplication using the steps in the following section:

After the conversion is complete, you can import the restored folder and add that folder as a path to the deduplication storage for Backup Exec.

Creating or importing deduplication disk storage

Deduplication disk storage provides a disk-based backup folder that you can use as a destination for backup jobs. When you use deduplication disk storage, only unique data is stored.

Before you create a deduplication disk storage device, review the requirements. It is recommended that you use a dedicated volume and a large amount of RAM for deduplication disk storage.

See [“Requirements for the Deduplication feature”](#) on page 972.

You can create only one deduplication disk storage device on a Backup Exec server. You can create deduplication disk storage on a storage array. However, if a deduplication disk storage device already exists on a Backup Exec server, then you cannot add another device to a storage array that is connected to that Backup Exec server.

When you use Backup Exec's **Delete** option on a deduplication disk storage device, the folder is removed from the Backup Exec database. However, the folder and the files in it remain on the disk. When you delete backup sets from a deduplication disk storage device, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

If you use the Backup Exec Central Admin Server feature, you can share a deduplication disk storage device between multiple Backup Exec servers. You can enable sharing when you add a deduplication disk storage device. You can

select new Backup Exec servers to share deduplication disk storage or remove the sharing ability for Backup Exec servers at any time.

You can create new deduplication disk storage, or you can import an existing deduplication disk storage device from another Backup Exec server.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 995.

Before you create or import a deduplication disk storage device, have the following information available:

- What you want to name the deduplication disk storage device.
- The volume on which you want to create the deduplication disk storage device.
 - The path of the existing deduplication disk storage device, if you import a deduplication disk storage device.
 - The information for the user account that was used when the existing deduplication disk storage device was originally created.
- The logon account to use to access the deduplication disk storage device. You cannot use the System Logon Account. It is recommended that you select or create a logon account that you use exclusively for the deduplication disk storage device. You should not use this account for any other purpose. This account should not contain credentials that are subject to password update policies.

The following requirements are for the password credentials for a deduplication disk storage device:

- The password cannot be blank
- The password cannot contain the following characters: &, ", <, >, %, ^
- The password cannot start with the hyphen character: -
- The password cannot end with the backslash character: \
- Whether you want to enable encryption while the data is transmitted to the deduplication disk storage device and while the data is stored on it. You should not use the Backup Exec encryption options for backup jobs that deduplicate data.
- The number of concurrent operations to run on the device. This setting determines the number of jobs that can run at the same time on this device. The number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. It is recommended that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.

To create or import deduplication disk storage

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server feature is not installed Select **Disk-based storage**, and then click **Next**.

If the Central Admin Server feature is installed Do the following in the order listed:

- Select the Backup Exec server that you want to configure deduplication disk storage for, and then click **Next**.
- Select **Disk-based storage**, and then click **Next**.

3 Click **Deduplication disk storage**, and then click **Next**.

4 Enter a name and description for the deduplication disk storage device, and then click **Next**.

5 Do one of the following:

To create a new deduplication disk storage device Click **Next** to accept the default location that appears in the field.

To import an existing deduplication disk storage device Enter the path of the existing deduplication disk storage device, and then click **Next**.

6 Click the drop-down arrow and select the logon account to use to access the deduplication disk storage device, or click **Add/Edit** to create a logon account, and then click **Next**.

7 Specify if you want to enable encryption during transmission of data to the deduplication disk storage device and while the data is stored on it, and then click **Next**.

8 Specify the number of concurrent operations that can run on the deduplication disk storage device, and then click **Next**.

9 Review the storage configuration summary, and do either of the following:

To change any of the selections

Do the following in the order listed:

- Click a review heading and make the appropriate changes.
- Click **Next** through the wizard to return to the summary screen.
- Click **Finish**.

To start the configuration

Click **Finish**.

Editing the properties of a deduplication disk storage device

You can edit some properties of a deduplication disk storage device.

To edit the properties of a deduplication disk storage device

- 1 On the **Storage** tab, double-click the name of the deduplication disk storage device.
- 2 In the left pane, select **Properties**.

3 Change the following properties as needed:

Name	Indicates the name that was entered when the deduplication disk storage was configured. You can change the name at any time.
Description	Indicates the description that was entered when the deduplication disk storage was configured. You can change the description at any time.
Logon account	Indicates the logon account that is being used to access the device. If you change the logon account after you configure the device, additional steps are required to enable access to the data that is already stored on the device. See “Changing the password for the logon account for deduplication disk storage” on page 990.
Encryption	Enables or disables encryption while the data is transmitted to the device, and while the data is stored on the device. You should not use the Backup Exec encryption options for backup jobs that deduplicate data.
Concurrent operations	Indicates the maximum number of jobs that you want to run at the same time on this device. The number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. It is recommended that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.
Data stream size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device being used.

Client-side deduplication	Indicates whether client-side deduplication is enabled for this device. Client-side deduplication enables a remote computer that is configured to send data directly to the deduplication disk storage. After the data is deduplicated, then only unique data is sent directly to the deduplication disk storage. By using this option, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations.
Percentage of disk space to reserve for non-Backup Exec operations	Displays the amount of disk space to set aside for applications other than Backup Exec. The default amount is 5%
Log level	Indicates the type of information you want to include in the diagnostic logs for this device. The choices range from critical errors only to all types of messages.
Log retention period	Indicates the number of days to keep the diagnostic logs for this device.

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the warning low disk space threshold.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 5%.

This property appears only if the deduplication disk storage is on a storage array.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the low disk space threshold.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 15%.

This property appears only if the deduplication disk storage is on a storage array.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 25%.

This property appears only if the deduplication disk storage is on a storage array.

- 4 Click **Apply** to save the changes

Changing the password for the logon account for deduplication disk storage

When you specify a Backup Exec logon account for a deduplication disk storage device, an additional user account is created for the deduplication components with the same user name and password. However, if you change the credentials for the Backup Exec logon account, the credentials for the additional user account are not changed automatically. You must use the `spauser.exe` utility to update the password for the additional user account. This account is known as the "User 1" account when you use the `spauser.exe` utility to view a list of user names that are associated with the deduplication disk storage.

To change the password for the logon account for deduplication disk storage

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account that you want to change, and then click **Edit**.

- 4 Type the current password for the logon account, and then click **OK**.
- 5 Click **Change Password**.
- 6 Type the new password in the **Password** field and in the **Confirm** field.
- 7 Click **OK**.
- 8 At a command prompt, switch to the Backup Exec program file directory, and then type the following command:

```
spausers.exe -c -u <UserName>
```

The default Backup Exec program file directory is C:<Backup Exec install path>\Backup Exec

The user name is case-sensitive. If you do not know the user name, type the following command to find the user name that is associated with "User 1":

```
spausers.exe -l
```

You are prompted for the old password and a new password. Ensure that the new password is the same as the password that you used in step 6.

Selecting storage devices for direct access sharing

Direct access enables a remote computer to send data directly to storage devices that are hosted by a Backup Exec server. When direct access sharing is enabled, the Backup Exec server is bypassed.

If you use a deduplication disk storage device or an OpenStorage device that supports client-side deduplication, then enabling direct access sharing enables Backup Exec to perform client-side deduplication. Note that client-side deduplication is CPU-intensive.

Direct access sharing becomes available after you create a backup job in which a deduplication device is selected and the following option is selected: **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported**.

After the correctly configured backup job is created, then the option **Direct access sharing** appears in the following locations.

- On the details screen for a server on the **Backup and Restore** tab.
- On the details screen for a storage device on the **Storage** tab.

In addition, the option **Direct access properties** appears on the details screen for a server on the **Backup and Restore** tab.

To select storage devices for direct access sharing

- 1 Do one of the following:
 - On the **Backup and Restore** tab, double-click the server that you want to set up to share devices.
 - On the **Storage** tab, double-click the storage device that you want to share.
- 2 In the left pane, select **Direct access sharing**.
- 3 Select the check box for the items that you want to share.

Note: After you enable direct access sharing for a client, you must restart the Backup Exec services on the Backup Exec server. Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**, and then click **Restart all services**.

See [“Editing server properties for direct access”](#) on page 992.

Editing server properties for direct access

For servers that are enabled for direct access, you can do the following:

- Add or change a description of the server.
- Enable or disable ICMP ping operations to detect the server.
- Add or edit a logon account that is used to access the remote computer.

To edit server properties for direct access

- 1 On the **Backup and Restore** tab, double-click the server that is enabled for direct access.
- 2 In the left pane, select **Direct access properties**.

3 Edit the following options as needed:

Server name	Indicates the name of the remote computer or managed Backup Exec server.
Description	Lets you enter a description of the server.
Port	Indicates the port that is used for communications between the Backup Exec server and the remote computer.
Use ICMP ping operations to detect the server	Lets the Backup Exec server use ICMP ping to locate the remote computer.
Logon account	Indicates the logon account that is required to access the remote computer. You can add a new logon account or edit an existing account.

4 To save changes, click **Apply**.

Changing the location of a deduplication disk storage device

You can change the location of an existing deduplication disk storage device. You can move the deduplication disk storage device to another volume on the same Backup Exec server. If the Central Admin Server feature is installed, you can move the deduplication disk storage device to another Backup Exec server. Only one deduplication disk storage device is supported per Backup Exec server.

It is recommended that you use the same name for the deduplication disk storage device when you change the location, but it is not required.

To change the location of a deduplication disk storage device

- 1 Ensure that no jobs are running or are scheduled to run until you have completed the process of changing the location of a deduplication disk storage device.
- 2 Document the current path of the deduplication disk storage device by viewing the properties of the device.
- 3 On the **Storage** tab, right-click the deduplication disk storage device, and then click **Disable**.
- 4 Right-click the deduplication disk storage device again, and then click **Delete**.

- 5 When you are prompted to delete the device, click **Yes**.
- 6 Click the Backup Exec button, select **Configuration and Settings**, select **Backup Exec Services**, and then click **Stop all services**.
- 7 In Windows Explorer, copy the deduplication disk storage device to the new path or volume.

Note: This step can take a long time as all of the data within the device is copied to the new location.

- 8 On the Backup Exec Administration Console, click the Backup Exec button, select **Configuration and Settings**, select **Backup Exec Services**, and then click **Restart all services**.
- 9 On the **Storage** tab, click **Configure Storage**, click **Disk-based storage**, and then click **Next**.
- 10 Click **Deduplication disk storage**, and then click **Next**.
- 11 Enter a name and description for the deduplication disk storage device, and then click **Next**.
- 12 Do one of the following:

If the deduplication disk storage device was created in Backup Exec 2012 and later	Click Create a new deduplication disk storage device .

If the deduplication disk storage device was created in a version prior to Backup Exec 2012	Click Import an existing deduplication disk storage device , and then enter the path to which you moved the deduplication disk storage device.

- 13 Click **Next**.
- 14 Specify the logon account that was used for the original deduplication disk storage device, and then click **Next**.
- 15 Specify if you want to enable encryption, and then click **Next**.
- 16 Specify the number of concurrent operations to run on the deduplication disk storage device, and then click **Next**.
- 17 Review the summary, and then click **Finish**.
- 18 When the deduplication disk storage device comes online, you can delete the original files.

Sharing a deduplication device between multiple Backup Exec servers

If you use the Backup Exec Central Admin Server feature, you can select which Backup Exec servers can share a deduplication disk storage device or an OpenStorage device. When you add a deduplication disk storage device or an OpenStorage device, the Backup Exec server that you used to add the device is automatically selected for sharing.

Note: To share a deduplication disk storage device, you must add it as an OpenStorage device on all Backup Exec servers that you want to access the device, except for the Backup Exec server that was used to create it.

This type of sharing is not the same as direct access sharing. With direct access sharing, a remote computer bypasses the Backup Exec server to directly access storage devices that are hosted by the Backup Exec server.

See [“Sharing storage devices”](#) on page 557.

See [“Selecting storage devices for direct access sharing”](#) on page 991.

How to use client-side deduplication

Client-side deduplication enables a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device. By using client-side deduplication, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If your deduplication device supports client-side deduplication, a remote computer deduplicates data and then sends only the unique data directly to a deduplication disk storage device or an OpenStorage device. Client-side deduplication is available for Windows computers and Linux computers.

Note: Client-side deduplication may increase the CPU utilization on the remote computer if your deduplication device supports client-side deduplication.

When you create a backup job with client-side deduplication, keep in mind the following items:

- The backup job can include resources from only one remote computer.
- The Agent for Windows is required on the remote Windows computer to perform Windows client-side deduplication. The Agent for Linux is required on the Linux computer to perform Linux client-side deduplication.

Note: A maximum of 64 remote agents with client-side deduplication enabled are allowed per Backup Exec server.

- The remote computer must be pingable.
- The remote computer cannot be a Backup Exec server.
- A deduplication disk storage device or an OpenStorage device must be used for the backup job.
- The option **Client-side deduplication** must be enabled on the properties for the storage device.
- The option **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported** must be selected in the **Storage** options for the backup job. This option is selected by default when you select a deduplication disk storage device or an OpenStorage device as the storage for a backup job.

If you do not configure the remote computer to use client-side deduplication, then the data from the remote computer is sent to the Backup Exec server to be deduplicated. Then, the deduplicated data is backed up to the deduplication disk storage or the OpenStorage device. This process increases the CPU utilization on the Backup Exec server. However, this process is useful if you are backing up older remote computers.

See [“About the Deduplication feature”](#) on page 969.

See [“How to set up backup jobs for deduplication”](#) on page 996.

See [“Editing the properties of an OpenStorage device”](#) on page 422.

See [“Editing the properties of a deduplication disk storage device”](#) on page 985.

How to set up backup jobs for deduplication

Set up a backup job for deduplication by selecting the option **Back Up to Deduplication Disk Storage**. Then, on the **Storage** settings, select either an OpenStorage device or a deduplication disk storage device as the destination device, and then select the deduplication method to use.

The following deduplication methods are available:

- If you want to enable client-side deduplication, select the option **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported**. This is the default option. If the storage device that you select for the job does not support client-side

deduplication, then either Backup Exec server-side deduplication or appliance deduplication is used.

- If you want to enable Backup Exec server-side deduplication, select the option **Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported**. If the storage device that you select for the job does not support server-side deduplication, then appliance deduplication is used.

See [“How to use client-side deduplication”](#) on page 995.

Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices

Backup Exec supports optimized duplication, which enables deduplicated data to be copied directly from one OpenStorage device to another OpenStorage device from the same vendor. Both devices must be attached to a single Backup Exec server. For example, you can copy data from one Arctera PureDisk device to another Arctera PureDisk device. Because the data is deduplicated, only unique data is copied between the devices.

To copy data between OpenStorage devices or deduplication disk storage devices, you must create a job to duplicate backup sets. The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job. No additional settings are required; optimized duplication occurs automatically when you set up a duplicate backup job between appropriate devices. You can restore data from either device.

Optimized duplication can be performed on backup sets that were enabled for Granular Recovery Technology (GRT). However, only deduplication disk storage devices and PureDisk devices support optimized duplication for GRT-enabled backup sets.

You can find a list of compatible types of storage devices in the Backup Exec Hardware Compatibility List.

Note: The OpenStorage devices must be from the same vendor. You cannot perform optimized duplication between OpenStorage devices from different vendors. If you attempt to copy deduplicated data between OpenStorage devices from different vendors, regular duplication is performed instead of optimized duplication.

Using optimized duplication with the Central Admin Server feature

If you use the Central Admin Server feature (CAS), the functionality of optimized duplication is expanded to let you do the following:

- Copy data from a deduplication disk storage device on one Backup Exec server to a deduplication disk storage device on another Backup Exec server.
- Copy data from an OpenStorage device that is attached to a Backup Exec server to another OpenStorage device that is attached to a different Backup Exec server.

To use optimized duplication with CAS, the following requirements must be met:

- You must have a license for the Enterprise Server feature. CAS is installed as part of the Enterprise Server feature.
- All Backup Exec servers that you use with CAS as either a central administration server or as managed Backup Exec servers must use the 64-bit version of Windows.
- You must have a central administration server and at least one managed Backup Exec server in your CAS environment.
- For client-side deduplication and Backup Exec server-side deduplication, you must configure one deduplication disk storage on the Backup Exec server from which you want to copy the deduplicated data. You must also configure one deduplication disk storage on the Backup Exec server to which you want to copy the deduplicated data.
- For appliance deduplication, the Backup Exec server from which you want to copy deduplicated data must have the appropriate plug-in for the OpenStorage device and a properly configured OpenStorage device. In addition, the Backup Exec server to which you want to copy the deduplication data must have the appropriate plug-in for the OpenStorage device and a properly configured OpenStorage device.
- You must share deduplication devices between the Backup Exec servers.
- You must inventory and catalog the media on the destination server before you recover any files from the duplicated backup set. You must do this regardless of how the catalog sharing option is configured for CAS.

Using optimized duplication to transfer backup data to a remote location

You can use optimized duplication to transfer backup data to a remote location over a WAN connection. You should prepopulate the destination deduplication disk storage device with a full backup of the servers. This prepopulation task is

also known as seeding the destination deduplication disk storage device. Seeding helps avoid the time-consuming and bandwidth-intensive process of transmitting large amounts of backup data over the low bandwidth WAN connection.

To seed a device, you can store a full backup to an external storage device such as a USB drive. You can then transport the USB drive to the remote location where the Backup Exec server and the destination deduplication disk storage device are kept, and duplicate the backup data to the deduplication disk storage device.

How to set up optimized duplication

Set up a duplicate backup job to perform optimized duplication.

Table B-5 How to set up optimized duplication

Step	For more information
If you are using CAS, do the following: ■ Verify that you have one central administration server and at least one managed Backup Exec server. ■ Verify that the Backup Exec server from which you want to copy the deduplicated data has a deduplication disk storage device (for client-side or Backup Exec server-side deduplication) or an OpenStorage device (for appliance deduplication). Also verify that the Backup Exec server to which you want to copy the deduplicated data has a deduplication disk storage device (for client or Backup Exec server-side deduplication) or an OpenStorage device (for appliance deduplication). ■ Verify that the Backup Exec servers are enabled for sharing. Note: This information applies only to CAS. If you do not have CAS, skip this step.	See “Configuring an OpenStorage device” on page 419. See “Creating or importing deduplication disk storage” on page 982.
Create a backup job that uses an OpenStorage device or a deduplication disk storage device as the destination.	See “Backing up data” on page 158.

Table B-5 How to set up optimized duplication (continued)

Step	For more information
Create a job to duplicate backup sets and select the appropriate OpenStorage device or deduplication disk storage as the destination. Note: The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job.	See “Duplicating backup sets or a job history manually” on page 222.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 995.

Copying deduplicated data to tapes

Backup Exec lets you copy deduplicated data from an OpenStorage device to tape for long-term or off-site storage. When data is copied to tape, it is rehydrated. In other words, the files are reassembled into their original form and are not deduplicated.

To copy deduplicated data to tapes, you must create a duplicate backup job that copies the backup sets from the OpenStorage device to a tape device.

See [“Duplicating backup sets or a job history manually”](#) on page 222.

Using deduplication with encryption

You should not use the Backup Exec encryption options for backup jobs that deduplicate data. Data cannot be deduplicated when the Backup Exec encryption options are used.

If you want deduplicated data to be encrypted on a deduplication disk storage device, you can enable the encryption property on the deduplication disk storage device.

See [“Encryption key management”](#) on page 733.

Restoring a deduplication disk storage device or deduplicated data

You can restore a deduplication disk storage by running the Restore Wizard. When you restore a deduplication disk storage device, the original folder is deleted and then replaced by the restored folder.

Note: You cannot redirect the restore of a deduplication disk storage device. You must restore a deduplication disk storage device to its original location.

To restore deduplicated data, you can create a regular restore job. No additional settings are required.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

To restore a deduplication disk storage device

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore the deduplication disk storage device, and then click **Restore**.
- 2 Select **Shadow Copy Components** and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Disaster recovery of deduplication disk storage devices”](#) on page 1001.

Disaster recovery of deduplication disk storage devices

A deduplication disk storage device is stored on the Backup Exec server. If your Backup Exec server experiences a disaster, then the data from the deduplication disk storage device is lost. Therefore, you should take steps to prepare for recovery from a system failure. To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication disk storage device. The snapshot includes the folder and the contents of the folder. You can store the snapshot on tape, which you can then use to recover your deduplication disk storage after you recover the Backup Exec server.

When you restore data from the snapshot, the following processes occur:

- Backup Exec stops the deduplication services if they are running. The deduplication services are separate from the Backup Exec services, so the Backup Exec services are not affected.
- Backup Exec deletes any files that are present in the deduplication disk storage.

- The deduplication disk storage is restored to its original location, along with the contents of the folder.
- The deduplication services are restarted.

Note: If you use Backup Exec Simplified Disaster Recovery (SDR) to recover the Backup Exec server, SDR does not recover the deduplication disk storage during the recovery of the Backup Exec server.

Preparing for disaster recovery of a deduplication disk storage device

To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication disk storage device. The snapshot includes the folder and the contents of the folder. You can store the snapshot on tape, which you can then use to recover your deduplication disk storage after a disaster.

To prepare for disaster recovery of a deduplication disk storage device

- 1 On the **Backup and Restore** tab, right-click the server where the deduplication disk storage device is located.
- 2 Select **Backup**, and then select **Backup to Tape**.
- 3 In the **Selections** box, click **Edit**.
- 4 Expand **Shadow Copy Components**, expand **User Data**, and then select **Backup Exec Deduplication Storage**.
- 5 Click **OK**.
- 6 Complete any additional options that you want to use.

It is recommended that you schedule this job to run just prior to the 12:20 a.m. and 12:20 p.m. deduplication maintenance times.

- 7 Click **OK** to create the job.

See [“Restoring a deduplication disk storage device or deduplicated data”](#) on page 1001.

Disaster recovery of OpenStorage devices

The following disaster recovery scenarios are possible for OpenStorage devices:

- The device fails.
- The Backup Exec server that uses the device fails.

If the device fails, you should consult the documentation from the device's vendor. If the Backup Exec server fails and you need to reinstall Backup Exec on the Backup Exec server, you must reconfigure the device, and inventory and catalog the media from it after the Backup Exec server is recovered.

Backup Exec Agent for VMware

This appendix includes the following topics:

- [About the Agent for VMware](#)
- [Requirements for using the Agent for VMware](#)
- [Using the Agent for VMware with Windows Server 2016 or later](#)
- [About installing the Agent for VMware](#)
- [Adding VMware vCenter Servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab](#)
- [Viewing details about VMware resources](#)
- [Installing the Agent for Windows on VMware virtual machines](#)
- [Push-installing the Agent for Windows to VMware virtual machines](#)
- [About establishing trust for a vCenter/ESX\(i\) server](#)
- [Backing up VMware virtual machines](#)
- [Restoring VMware virtual machines and vmdk files](#)
- [About instant recovery of a VMware virtual machine](#)
- [Troubleshooting the Agent for VMware](#)
- [About Recovery Ready for VMware virtual machines](#)

About the Agent for VMware

The Backup Exec Agent for VMware (Agent for VMware) lets you back up and restore virtual machines that use VMware ESX/ESXi/vSphere/vCenter.

You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

Backup Exec performs a single-pass backup to protect all guest virtual machines and the VSS-aware applications that are installed on the guest virtual machines. Backup Exec's Granular Recovery Technology (GRT) is enabled by default for files and folders. You can use a GRT-enabled backup to restore individual files and folders from a Windows guest virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange, SQL, SharePoint, and Active Directory applications that reside on guest virtual machines if you select the options for application-level GRT in your backup jobs. Note that application-level GRT is not enabled by default. You must select the application-level GRT options that you want to use when you set up a backup job.

For information about the best practices to use Backup Exec Agent for VMware, refer to *Backup Exec Best Practices*.

Additional features of the Agent for VMware let you do the following:

- Redirect the restore of data from a guest virtual machine to an alternate folder, datastore, host, or network.
- Back up to a disk device or to a tape device.
- Perform incremental and differential backup jobs. This feature is available only if your virtual machines are configured with hardware version 7 or later.

See [“Requirements for using the Agent for VMware”](#) on page 1005.

See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 1028.

See [“Backing up VMware virtual machines”](#) on page 1013.

See [“Restoring VMware virtual machines and vmdk files”](#) on page 1034.

Requirements for using the Agent for VMware

Before you use the Agent for VMware, ensure that the following requirements are met:

- Install a license for the Agent for Hyper-V and VMware on the Backup Exec server. The Backup Exec server runs the backup and restore jobs.

- Determine if you are going to use Backup Exec's Granular Recovery Technology to recover individual items from Microsoft applications. If you plan to use GRT, install the Agent for Windows on your virtual machines that run Windows.

Note: If you are going to use GRT, ensure that you use unique virtual machine names. GRT operations may not function correctly if duplicate virtual machine names are encountered.

See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 1028.

See [“Installing the Agent for Windows on VMware virtual machines”](#) on page 1010.

- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List.
You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.
- Verify that your virtual machine names do not contain any characters that VMware does not support. Only the following characters can be used in virtual machine names:
 - Uppercase and lowercase ASCII characters
 - Numbers
 - Period (.)
 - Hyphen (-)
 - Underscore (_)
 - Plus sign (+)
 - Left and right parentheses ()
 - Spaces
Use of any characters that are not listed above may cause backup jobs to fail.
- Verify that HTTPS port 443 is used on the vCenter/ESXi server. Port 443 is the default HTTPS port. Backup Exec is configured to use HTTPS port 443 to retrieve the appropriate SSL certificate from the vCenter/ESXi server. If your vCenter server does not use the default HTTPS port of 443, then you must add the server to Backup Exec with the port number appended to the end of the server name.

For example, myvCenter60.mydomain:482

- If you use VMware vCenter Server 6, verify that you do not have a mixed IPv4 and IPv6 configuration. VMware vCenter Server 6.0 does not support mixed IPv4 and IPv6 configurations. Backup and restore jobs may fail in this configuration.
- To use the new CreateSnapshotEx API in Backup Exec 25.0 or later, ensure that the following requirements are met:
 - vSphere 6.5 or later is installed.
 - VMware Tools 10.1.0 or later must be installed on the guest virtual machine. (VMware Tools 10.1.0 is released with ESXi 6.5, and earlier versions are at end of general support).
 - By default, VMware VSS provider is required to take the quiesce snapshots. If the VMware VSS provider is not running, quiesce snapshot will fail. For details, see <https://kb.vmware.com/s/article/82544>.

Using the Agent for VMware with Windows Server 2016 or later

The following information applies to virtual machines that use Windows Server 2016 or later:

- If the Backup Exec server is running an operating system prior to Windows 2016, Granular Recovery Technology (both application-level and file/folder level GRT) cannot be performed for a virtual machine that has a Resilient File System (ReFS) volume. A backup job for a virtual machine that is configured with Windows ReFS and is enabled for GRT will complete with a status of success with exceptions.

Note: Backup Exec supports file/folder GRT for ReFS volumes or application-level GRT for a virtual machine only if the Backup Exec server runs the same or a higher version of the operating system than the virtual machine.

For example, if the virtual machine runs on Windows 2016 and has REFS volumes, then the Backup Exec server should also run on Windows 2016 or later.

- If the Backup Exec server is running an operating system prior to Windows 2016, Granular Recovery Technology (both application-level and file/folder level GRT) cannot be performed for the virtual machine that contains volumes

on which Windows deduplication is enabled. A backup job for a virtual machine that is configured with a Windows deduplication volume and is enabled for GRT will complete with a status of success with exceptions.

See [“About the Agent for VMware”](#) on page 1005.

About installing the Agent for VMware

The Agent for VMware is installed as part of the Agent for Hyper-V and VMware. Install this license on the Backup Exec server to enable the Agent for VMware. You do not have to install the agent on the host server.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

Adding VMware vCenter Servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab

You can add VMware vCenter Servers and ESX/ESXi hosts to the list of servers on the **Backup and Restore** tab so that you can back up the virtual machines that they host.

When you add these hosts, the Backup Exec UI displays the "Trust not established" message if a trust is not established between the hosts and the media server. However it does not fail the backup or restore. You can verify the certificate and establish the trust from the Backup Exec UI.

To add a VMware vCenter Server or ESX/ESXi host to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **VMware vCenter server or ESX host**, and then click **Next**.
- 3 Select **Add a VMware vCenter server or ESX host to the list of servers**.
- 4 Enter the name or IP address of the server you want to add, and then add an optional description of the server.

- 5 If you want Backup Exec to install the Agent for Windows on any guest virtual machines for this host server, verify that the option **After adding the virtual host, install the Backup Exec Agent for Windows on the guest virtual machines** is selected.

To perform Granular Recovery Technology (GRT)-enabled backups of the Microsoft applications that are installed on the virtual machines, the Agent for Windows must be installed on the guest virtual machines. If you do not want to install the Agent for Windows on the guest virtual machines, clear the check box. You can install the Agent for Windows later if you decide that you want to use GRT.

- 6 Click **Next**.
- 7 Select the correct logon account for the server from the drop-down list.

If the logon account for the server is not in the list, click **Add/Edit** to add it to the list.
- 8 Click **Next**.
- 9 Review the summary information, and then click **Finish**.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

See [“Installing the Agent for Windows on VMware virtual machines”](#) on page 1010.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 953.

Viewing details about VMware resources

The **Backup and Restore** tab includes a **Details** pane on the lower portion of the screen. The **Details** pane provides additional details for any type of server that is selected from the list of servers in the upper portion of the screen. Additional details and functionality appear if a VMware host is selected. The **Details** pane for VMware resources includes details about the last 7 days of backup jobs, the date of the last backup, and the date of the next scheduled backup. In addition, it includes the ability to back up and restore data, and to filter the list of guest virtual machines.

The **Details** pane for VMware virtual machines includes the resources that may not appear in the list of servers, such as:

- Virtual machines that do not have the Agent for Windows installed.
- Non-Windows virtual machines, such as Linux.
- Virtual machines that do not have a GRT-enabled backup.

The **Details** pane for VMware virtual machines provides a combination of current information and historical information. The **Refresh** button enables Backup Exec to discover all server resources. However, if a virtual machine has been moved, deleted, or has had a credentials change, then the following occurs:

- If the virtual machine has never been backed up, Backup Exec removes it from the **Details** pane.
- If the virtual machine has been backed up, Backup Exec does not remove it from the **Details** pane.

If a virtual machine is renamed, then the following occurs in the **Details** pane:

- If the virtual machine has been backed up, the **Details** pane includes an entry for the new name and keeps the entry for the old name.
- If the virtual machine has never been backed up, the **Details** pane includes only the new virtual machine name.

To view details about VMware resources

- ◆ On the **Backup and Restore** tab, select a VMware host from the list of servers.
The resources for the selected host appear in the **Details** pane.

Installing the Agent for Windows on VMware virtual machines

To use Backup Exec's Granular Recovery Technology (GRT) for Microsoft application data, install the Agent for Windows on any virtual machines that run Windows. To perform this procedure, you must have already added the vCenter or ESX/ESXi host to the list of servers on the **Backup and Restore** tab.

Note: VMware Tools should be installed before you install the Agent for Windows on a guest virtual machine if you intend to use the VMware VSS provider.

See [“Adding VMware vCenter Servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab”](#) on page 1008.

To install the Agent for Windows on VMware virtual machines

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **VMware vCenter or ESX host**, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the virtual machines of a VMware vCenter server or ESX host**.

- 4 Select the server from the drop-down list, and then click **Next**.
- 5 Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 6 Select the guest virtual machines that you want to install the Agent for Windows on, and then click **Next**.
- 7 Select the logon account for the guest virtual machines, and then click **Next**.
- 8 Select the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically

Select this option to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.

Note: If you simply want to reestablish the trust with the existing version of the Agent for Windows that is already installed on the selected virtual machines, you should uncheck this check box.

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required

Select this option to enable Backup Exec to automatically restart the remote computer, if required.

- 9 Review the summary, and then click **Install**.

Push-installing the Agent for Windows to VMware virtual machines

You use the **Add Server** wizard to push-install the Agent for Windows to the VMware virtual machines that you select. You can choose to install the Agent for Windows on all of the guest virtual machines that need it or on specific machines only. For example, if you know that some of your guest virtual machines are scratch machines and do not need to be backed up, you can exclude those virtual machines from the installation.

Note: VMware Tools should be installed before you install the Agent for Windows on a guest virtual machine if you intend to use the VMware VSS provider.

To push-install the Agent for Windows to virtual machines

- 1 On the **Backup and Restore** tab, right-click the virtual host, and then select **Install Backup Exec Agent for Windows to guest virtual machines**.
- 2 Check the **Allow Backup Exec to establish a trust with the servers** option, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the guest virtual machines of a VMware vCenter or ESX server**.
- 4 In the **VMware vCenter or ESX server** field, select the VMware server that hosts the virtual machine.
- 5 Click **Next**.
- 6 Select the guest virtual machines on which you want to install the Agent for Windows, and then click **Next**.
- 7 Verify that the virtual machines you selected are online and select the appropriate logon account for those virtual machines, and then click **Next**.
- 8 Review the information on the **Summary** panel, and then click **Install**.

About establishing trust for a vCenter/ESX(i) server

When you add the hosts, the Backup Exec Media Server provides an option to establish trust with the virtual host server. As part of the Establish Trust workflow, Backup Exec attempts to validate the SSL certificate presented by the host.

A commonly used third party CA-signed certificate is automatically validated.

If the certificate is self-signed, the certificate cannot be recognized. You must validate the certificate for the trust to be established. Backup Exec saves required information of the SSL certificate of the vCenter Server or ESX(i) host into the Backup Exec database. You do not have to validate the certificate again on subsequent connections to the server.

Backup Exec automatically attempts to validate the certificates for all virtual hosts on a scheduled daily basis. If there is any change in the server certificate, it is detected and the server is marked as an untrusted server. A backup or restore job does not fail if the job is run for a server that is marked as untrusted. If the validity of the certificate expires, the server is marked as untrusted.

On the Backup Exec console the status displays **Trust not Established** for an untrusted server.

Establishing trust for a vCenter/ESX(i) server

You can establish trust for a vCenter/ESX(i) server from the **Backup and Restore** tab.

To establish trust for a vCenter/ESX(i) server

- 1 On the **Backup and Restore** tab, right-click the server for which you want to establish trust and then click **Establish Trust**.
- 2 If the certificate is not automatically recognized, on the **Establish Trust** dialog box, click **View Certificate** to validate the identity of the host.
- 3 If the certificate is valid, click **Yes** to establish trust.

If trust is established on the Central Administration Server (CAS), the status is automatically passed on to all the Managed Media Servers (MMS).

Backing up VMware virtual machines

When you create a backup job, you can select an entire vCenter server or ESX/ESXi host, datacenters, folders, or individual virtual machines. Additionally, Backup Exec's dynamic inclusion feature can automatically back up new virtual machines and folders that are found when a backup job runs. If you select the vCenter server or the ESX/ESXi host for a backup job, all virtual machines on that host are selected for backup automatically. However, you can edit the selections to include only selected virtual machines in the backup job. A backup of a vCenter server or ESX/ESXi host does not include independent disks or configuration files for the vCenter server or ESX/ESXi host.

Notes about backing up VMware virtual machines

You should review the following notes before you create backup jobs for VMware virtual machines:

- If you attempt to back up a virtual machine with the NetBIOS name "LocalHost", the backup will fail with the error "e000976f - Cannot backup the virtual machine to the deduplication device that is hosted by the same virtual machine."
- Physical and Virtual Raw Disk Mapping (RDM) disks are not backed up. All other disks are backed up.
- If you want to back up an Exchange database availability group (DAG) on a virtual machine, use the Backup Exec Agent for Exchange. The Agent for VMware does not support backups of Exchange DAG.
- You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

- If you back up a virtual machine that runs any supported version of Microsoft SharePoint, you may experience a processing time of more than 30 minutes.

To back up VMware virtual machines

- 1 On the **Backup and Restore** tab, select a virtual machine, an ESXi host, or a vCenter server from the list of servers.
- 2 Do one of the following:

If you selected an ESXi host or a vCenter server in step 1

In the **Backups** group at the top of the screen, click **Backup**, and then select the type of backup you want to perform.

Alternatively, in the **Details** pane at the bottom of the screen, click **Backup**, and then select the type of backup you want to perform.

If you selected a virtual machine in step 1

Do the following:

- In the **Backups** group at the top of the screen, click **Backup**, and then select the type of backup you want to perform.
- On the **Back Up Virtual Machine** panel, select the backup method that you want to use for this virtual machine, either **Virtual-based backup** or **Agent-based backup**, and then click **Next**.

For information about which backup method to use, see the following topic:

See [“Recommendations for when to use virtual-based backup and agent-based backup”](#) on page 139.

Note: If you selected virtual-based backup and the Agent for VMware and Hyper-V is not installed, you will be prompted to either install it or to create an agent-based backup instead.

- If you selected **Virtual-based backup**, select the virtual machine's host, and then click **Next**.

If the virtual machine's host is not in the list of servers, click **Add**, and then complete the steps in the wizard to add the host.

- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit** to add or remove resources from the backup selection list.
- 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.
- 5 Select the disks for a virtual machine that you want to back up.

You can select the entire virtual machine (select all disks) or select disks individually. You can also exclude disks selectively.
- 6 Click **OK**.

If you make partial selections, the **Virtual machines have partially selected disks** pop-up is displayed.

The job selections consist of one or more virtual machines where only some of its virtual disks are selected. For a virtual machine to function at the operating system level and application level, ensure that all the required disks are selected.

To perform File/Folder GRT and application GRT restore, ensure that the system disk is selected. For application GRT, select all virtual disks with application data.
- 7 Click **OK**.
- 8 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 9 On the **Backup Options** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.
- 10 On the **Backup Options** dialog box, in the left pane, select **Virtual Machines**.
- 11 Set any of the following options for this job:

Item	Description
Maximum number of streams used for parallel backup	Select the maximum number of streams you want to use for parallel backup of virtual disks of a virtual machine. You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage. Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups. See “Multi-stream backup of virtual machines” on page 1033.
Minimum number of streams used for parallel backup	Select the minimum number of streams you want to use for parallel backup of virtual disks of a virtual machine. You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage. The concurrent write sessions value on the disk storage device must be greater than the minimum number of streams that are used for parallel backups, or the backup job fails. Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups. See “Multi-stream backup of virtual machines” on page 1033.

Item	Description
Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option if you want Backup Exec to perform a full backup if an incremental backup or a differential backup cannot be performed. If you do not select this option and Backup Exec cannot perform an incremental backup or a differential backup, then the job fails. In addition, if Backup Exec detects a configuration change, then a full backup must be performed. If a configuration change is detected and Backup Exec cannot perform a full backup, then the job fails if this option is not selected. This scenario applies only if a full backup and some incremental backups or differential backups have already been performed and the next scheduled job is for an incremental backup or a differential backup.
Back up virtual machines that are powered off	Select this option if you want Backup Exec to back up any virtual machines that are turned off.
Enable Instant Recovery for all virtual machines, even those not eligible for GRT	Select this option to enable Instant Recovery of all virtual machines, even those that run operating systems which are not eligible for Granular Restore Technology.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the backup. This option is for the virtual machines that use a Windows operating system only. The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk. Note: GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers. See “Using Granular Recovery Technology (GRT) with the Agent for VMware” on page 1028.
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Active Directory is installed.

Item	Description
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Exchange is installed.
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SQL is installed.
Run a SQL Log backup after backing up the virtual machine	Select this option to enable Backup Exec to back up the SQL logs for the databases that use logging. After the logs are backed up, the data from the logs is committed to the database and the log is emptied so that it can receive new data. If this option is not selected, the SQL logs continue to grow until the disk is full or until you perform a manual backup job to back up the logs.
Enable GRT for Microsoft SharePoint on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SharePoint items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SharePoint is installed.
Back up using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	Select this option to enable Backup Exec to create a quiesced snapshot for the virtual machine. If the quiesced snapshot cannot be created, then Backup Exec creates a crash-consistent snapshot instead. Note: Backup Exec automatically selects this option if you select any of the options for enabling GRT for Microsoft applications.

Item	Description
Use VSS Copy backup	Select this option to enable Backup Exec to perform a VSS copy backup instead of a full backup. If you do not select this option (by default not selected), Backup Exec performs a full VSS backup. The VSS Provider initiates snapshots using the Full backup type setting. Each application responds differently to this request. For example, iIn the case of Microsoft Exchange, the database logs are truncated if this option is not selected.

Item	Description
Transport mode priority list	Select the method to transport the Virtual Machine Disk Format (vmdk) file from the ESX/ESXi host. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority. The following methods are available: ■ SAN - Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the Backup Exec server can access. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi host is not affected. ■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist: ■ The virtual machine is not located on the SAN. ■ The Backup Exec server does not have access to the SAN. ■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist: ■ The virtual machine is not located on the SAN. ■ The Backup Exec server does not have access to the SAN. ■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hotadd feature for ESX/ESXi. See your VMware documentation for more information about hotadd. The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.

Item	Description
Enable asynchronous read for improved performance	Select this option to enable Backup Exec to perform asynchronous read for backup jobs on the source ESXi host. Asynchronous read is only available for NBD and NBDSSL transport modes and from vSphere 6.7 or later. ■ Number of read requests: Specify the number of active asynchronous read requests to the ESXi host. The default number is 8. ■ Buffer size: Specify the amount of data that you want to read from one asynchronous read request. The default size is 1 MB. For more information about asynchronous read, refer to the VMware documentation.
Backup method	Select the backup method that you want to use for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties. See “Configuring backup methods for backup jobs” on page 197. Note: Backup method is not applicable for Forever Incremental backup jobs.

- 12 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
- 13 Click **OK**.
- 14 On the **Backup Definition Properties** dialog box, click **OK**.

If you selected the virtual-based backup method, the backup job appears in the jobs list under the virtual host's name or IP address

See [“How Backup Exec automatically backs up new VMware virtual machines during a backup job ”](#) on page 1027.

See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 1028.

Setting default backup options for virtual machines

You can use the defaults that Backup Exec sets during installation for all VMware backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for virtual machines

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for virtual machine backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select **Virtual Machines**.
- 4 Select the appropriate options.

Item	Description
Maximum number of streams used for parallel backup	Select the maximum number of streams you want to use for parallel backup of virtual disks of a virtual machine. You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage. Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups. See "Multi-stream backup of virtual machines" on page 1033.

Item	Description
Minimum number of streams used for parallel backup	Select the minimum number of streams you want to use for parallel backup of virtual disks of a virtual machine. You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage. The concurrent write sessions value on the disk storage device must be greater than the minimum number of streams that are used for parallel backups, or the backup job fails. Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups. See “Multi-stream backup of virtual machines” on page 1033.
Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to perform a full backup if an incremental backup or a differential backup cannot be performed. If you do not select this option and Backup Exec cannot perform an incremental backup or a differential backup, then the job fails. In addition, if Backup Exec detects a configuration change, then a full backup must be performed. If a configuration change is detected and Backup Exec cannot perform a full backup, then the job fails if this option is not selected. This scenario applies only if a full backup and some incremental backups or differential backups have already been performed and the next scheduled job is for an incremental backup or a differential backup.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines that are turned off.
Enable Instant Recovery for all virtual machines, even those not eligible for GRT	Select this option to enable Instant Recovery of all virtual machines, even those that run operating systems which are not eligible for Granular Restore Technology.

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the backup. This option is for virtual machines that use a Windows operating system only. The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk. Note: GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers.
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Exchange is installed.
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SQL is installed.
Run a SQL Log backup after backing up the virtual machine	Select this option to enable Backup Exec to back up the SQL logs for databases that use logging. After the logs are backed up, the data from the logs is committed to the database and the log is emptied so that it can receive new data. If this option is not selected, the SQL logs continue to grow until the disk is full or until you perform a manual backup job to back up the logs.

Item	Description
Enable GRT for Microsoft SharePoint on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SharePoint items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SharePoint is installed.
Exclude virtual machines that must be put into a saved state for backup	Select this option to exclude from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins. This option applies only to the Agent for Hyper-V.
Back up using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	Select this option to enable Backup Exec to create a quiesced snapshot for the virtual machine. If the quiesced snapshot cannot be created, then Backup Exec creates a crash-consistent snapshot instead. Note: Backup Exec automatically selects this option if you select any of the options for enabling GRT for Microsoft applications.
Use VSS Copy backup	Select this option to enable Backup Exec to perform a VSS copy backup instead of a full backup. By default the VSS Provider initiates snapshots using the Full backup type setting. Each application responds differently to this request. In the case of Microsoft Exchange, the database logs are truncated. If you do not select this option, Backup Exec performs a full VSS backup

Item	Description
Transport mode priority list	Select the method to transport the Virtual Machine Disk Format (vmdk) file from the ESX/ESXi host. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority. The following methods are available: ■ SAN - Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the Backup Exec server can access. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi host is not affected. ■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist: ■ The virtual machine is not located on the SAN. ■ The Backup Exec server does not have access to the SAN. ■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist: ■ The virtual machine is not located on the SAN. ■ The Backup Exec server does not have access to the SAN. ■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hotadd feature for ESX/ESXi. See your VMware documentation for more information about hotadd. The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.

Item	Description
Enable asynchronous read for improved performance	Select this option to enable Backup Exec to perform asynchronous read for backup jobs on the source ESXi host. Asynchronous read is only available for NBD and NBDSSL transport modes and from vSphere 6.7 or later. ■ Number of read requests: Specify the number of active asynchronous read requests to the ESXi host. The default number is 8. ■ Buffer size: Specify the amount of data that you want to read from one asynchronous read request. The default size is 1 MB. For more information about asynchronous read, refer to the VMware documentation.
Backup method	Select the backup method that you want to use for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties. Note: Backup method is not applicable for Forever Incremental backup jobs. See “Configuring backup methods for backup jobs” on page 197.

5 Click **OK**.

See [“Backing up VMware virtual machines”](#) on page 1013.

How Backup Exec automatically backs up new VMware virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

If you select a VMware server, then dynamic inclusion is enabled automatically for all of the nodes below it that have a folder icon. If no virtual machines are located during a backup job, then the job fails.

See [“Backing up VMware virtual machines”](#) on page 1013.

Using Granular Recovery Technology (GRT) with the Agent for VMware

Backup Exec's Granular Recovery Technology (GRT) lets you restore individual drives, files, and folders from backup sets without having to restore the entire virtual machine. It also lets you restore individual items from the following VSS-aware applications that reside on virtual machines:

Table C-1 Types of data that Backup Exec backs up for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec backs up
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units
Microsoft SharePoint	SharePoint data

GRT works only for the virtual machines that use a Windows operating system. GRT does not work for system recovery.

GRT must be enabled in the VMware backup job. When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using GRT. By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

During the backup job, Backup Exec collects metadata from the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

Requirements for using GRT to back up Microsoft application data on virtual machines

The following items are required to back up data for Microsoft Exchange, SQL, SharePoint, and Active Directory on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.
- The Backup Exec server must be able to connect to the virtual machine using the network name or IP address.
- VMware Tools should be installed on the virtual machine.
- The Backup Exec Agent for Windows must be installed on the virtual machine. Be sure that the VMware Tools are installed before you install the Agent for Windows.

Note: If you installed the VMware Tools after you installed the Agent for Windows, you should copy the file `freeze.bat` from Backup Exec RAWs Install Path\VSS Provider to VMware Tools install path\backscripts.d, for example `\\Program Files\VMware\VMware Tools\backscripts.d`. If the `backscripts.d` directory does not exist, you must create it manually. Alternatively, instead of moving the `freeze.bat` file, you can repair the Agent for Windows, which automatically places all missing files in the appropriate location.

- The Backup Exec Agent for Applications and Databases must be installed on the Backup Exec server.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.
- The options for application-level GRT must be selected in backup jobs. These options are not selected by default.

Unsupported configurations for GRT

Before you create a GRT-enabled backup job for VMware resources, review the following information to understand what configurations are not supported for GRT.

Table C-2 Unsupported configurations for GRT

Unsupported items	Details
Virtual machines that have a combination of independent and non-independent disks	Backup Exec does not support GRT of independent disks.
Virtual machines that have a virtual RDM disk	If you try to restore a backup of a virtual machine that has a virtual RDM disk, the vmdk that corresponds to the virtual RDM disk cannot be created or restored. The restore job fails with the error "Unable to open a disk of the virtual machine". Only non-virtual RDM disks can be restored.
Virtual machines that have RAID 5 volumes	Backup Exec does not support file/folder-GRT of RAID 5 volumes. Application-level GRT is also not supported for virtual machines if one of the volumes on the virtual machine is a RAID 5 volume.
Virtual machines that have NTFS with unnamed mount points	Backup Exec does not support file/folder-GRT of NTFS with unnamed mount points. The restore job fails with the error "Unable to attach to a resource. Ensure that the selected resource exists and is online, and then try again." Backup Exec does not support Application GRT for VMware when the application is on an unnamed mount point.
Virtual machines that have utility partitions	File/folder-level GRT restore fails with the error "Unable to attach to a resource. Ensure that the selected resource exists and is online, and then try again." Backup Exec does not support backups of virtual machines that have utility partitions. Note that this is not the same as unnamed partitions.
Restores of full and incremental backup sets from different storage devices	Backup Exec does not support restores from mixed media if GRT was enabled in the backup job. For example, if the full backup is on tape and the incremental backup is on a disk storage device, the restore job will fail. Restores from mixed media types are supported if GRT is not enabled.

Table C-2 Unsupported configurations for GRT (*continued*)

Unsupported items	Details
Virtual machines that have dynamic disks (with GPT partition style)	Backup Exec does not support granular recovery of files, folders, and applications on virtual machines that have dynamic disks (with GPT partition style).
Virtual machines that have ReFS and Deduplicated volumes	Backup Exec does not support file/folder GRT for ReFS and Deduplicated volumes or application-level GRT for virtual machines if the Backup Exec server is not running Windows 2012 or later.

See [“Backing up VMware virtual machines”](#) on page 1013.

How cataloging works with VMware virtual machine backups

When you enable Granular Recovery Technology (GRT) for a backup job of a virtual machine, you can choose to run the catalog job for GRT as part of the backup job, as a separate job immediately after the backup job completes, or according to a schedule. By default, the catalog operation runs immediately after the backup job completes.

Note: The Instant GRT or full catalog options are not supported for backups to tape.

The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. You may want to schedule the catalog operation to run outside of your backup window so that it does not interfere with backup jobs. If the catalog operation is scheduled, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery on VMware virtual machines. Before the full catalog job completes, instead of using the Search wizard, you must browse the backup sets to select the items that you want to restore.

For example, if you set up incremental backups to run every 11 hours and set up the catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)
- Incremental 1 (10:00 P.M.)

- Catalog 1 (Midnight). This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)
- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight). This job catalogs Incremental 5. Incremental 4 is not cataloged.
- Incremental 6 (5:00 A.M.). This backup is not cataloged.

In the example, the full catalog operation runs only for Incremental 5, Incremental 3, and Incremental 1. For such jobs, you can use the Search wizard to search the data or you can quickly browse for individual items that you want to restore. You can perform a granular recovery using Incremental 2, Incremental 4, and Incremental 6 as well; however, it takes slightly longer to browse items because they are not fully cataloged. Backup Exec dynamically displays the granular data by mounting the backup set.

How byte count is calculated for Instant GRT or full catalog operations

On the **Job Monitor** and **Job History**, the byte count that displays for a catalog operation may differ from the byte count that displays for the corresponding backup job. The byte count for a catalog job may be larger than the byte count for a backup job. The way in which Backup Exec catalogs the data affects the byte count that appears for the catalog job.

- When a catalog operation is performed for a full backup, the data is read on a file-by-file basis and the byte count is calculated accordingly. During the full backup job, the data is read in terms of the number of sectors and the byte count is calculated based on the number of sectors. Therefore, the byte count for the catalog job may be larger than the byte count for the backup job.
- When the catalog operation is performed for an incremental backup, all files on the virtual disk are cataloged instead of only the changed files. Therefore, the byte count for the catalog job takes into account both the full backup and the incremental backup.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Multi-stream backup of virtual machines

Backup Exec now supports multi-stream backup of disks within a virtual machine. Virtual disks are backed up in parallel, which helps to achieve faster virtual machine backups.

This feature provides the following advantages:

- Faster backups
- Optimized backup window
- Enhanced utilization of disk performance and network bandwidth

Backup Exec provides options to change the backup streams as required. Each stream performs the backup of a virtual disk concurrently. You can configure the minimum number and maximum number of streams to be used for the backup of a virtual machine.

See [“Backing up VMware virtual machines”](#) on page 1013.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 1074.

You also have the flexibility to control the number of streams that can run in parallel across all jobs configured to back up different virtual machines. You can do this with the number of write sessions available in the target storage configuration. Ensure that the number of write sessions are configured based on the number of streams that are configured across jobs, which are scheduled to run concurrently.

Backup Exec Remote agent (RAWS) must be updated on the Hyper-V host to use multi-stream for Hyper-V virtual machines.

The following factors are important and affect backup performance when you use multiple streams:

- Available network bandwidth
- Source and target disk read-write throughput
- Number of parallel streams
- Locations of virtual disks on different source storage

Ensure that these factors are considered when you set the maximum number and minimum count of streams. If throttling is observed for any of the factors, it can affect the backup performance. You must have adequate throughput across all resources such as, source and target disk, CPU, and network.

This feature supports multiple streams only for the backup operation. The verify, duplicate, and restore operations are not supported and they continue to work with a single stream.

You can configure multiple streams only for backup-to-disk and deduplication disk storage.

The following target storage devices are not supported for multi-stream backups:

- Tape
- Cloud

For these devices, a single stream is used for backup.

Notes for using parallel streams

- If a virtual machine has only one disk, Backup Exec uses only one stream.
- Multi-stream is not enabled by default for the backup jobs. You must change the stream count to enable multi-stream backups.
- If the number of write sessions is less than 2, multi-streams backups are not enabled, and all jobs run as a single stream backup.

Restoring VMware virtual machines and vmdk files

You can use the Restore Wizard to restore the following:

- A complete virtual machine.
- The Virtual Machine Disk Format (vmdk) file for a virtual machine.
- Individual files and folders that were backed up from inside the vmdk file and individual items from Microsoft SharePoint, Exchange, SQL, or Active Directory. The ability to restore individual files and folders is available only if Granular Recovery Technology (GRT) was selected for the backup job.

Note: GRT cannot restore system state files such as the active registry.

Backup Exec lets you restore VMware data to either the location from which the data was backed up or to a different location. Restoring data to a different location is referred to as a redirected restore.

A redirected restore is useful for disaster recovery situations. When you perform a redirected restore, you have the option to restore the virtual machine to the most recent hardware version that the destination environment supports. If you do not select the option to restore the virtual machine to the most recent hardware version, then the virtual machine's original hardware version is preserved when it is restored.

For restore, if you select disks from within a virtual machine backup, instead of deleting the entire virtual machine, only the selected disks are restored and the data from the remaining disks is retained. Ensure that the disks required for a virtual machine to work at an operating system level and application level are selected while doing a restore. If you are doing a redirected restore where partial disks are selected and the target datastore is same as the original datastore, the Virtual machine name must be different than the original name. During partial restore, if snapshots exist for the virtual machine, you must do a full virtual machine restore.

Note: The steps below apply to virtual machines that were backed up with the virtual-based backup method. If you backed up the virtual machine using the agent-based backup method, follow the steps for restoring a non-virtual backup.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 241.

See [“To restore VMware virtual machines or vmdk files to the same location from which they were backed up”](#) on page 1035.

See [“To redirect the restore of VMware virtual machines or vmdk files to a different host”](#) on page 1038.

See [“To restore to a different path”](#) on page 1042.

To restore VMware virtual machines or vmdk files to the same location from which they were backed up

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled data**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
 - Click **Restore**, and then select **Restore virtual machine from the host**.
 - In the **Restore Wizard**, select **VMware data**, and then click **Next**.
- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
 - 3 If the **Where do you want to restore the data from** panel appears, verify that you want to restore from the selected storage or select different storage to restore from, and then click **Next**.
 - 4 Select **To the original location**.
 - 5 Select **Scan backup sets for malware before restore** to scan the selected backup set for malware before restore is performed.

For an incremental backup, the previous backup sets in the recovery chain are also used. To restore backup sets even if the malware scan fails, select the **Continue to restore even if scan fails** check box.

Malware scan operation is CPU intensive and can take significant time if the virtual machine backup has a large amount of data. In this scenario the time that the restore job requires can increase and you can schedule the scan. Scan operation is only supported for GRT enabled sets on disk and deduplication devices.

For GRT restore, if you are running a malware scan on backup sets as part of the restore job or ad hoc scan, the entire virtual machine backup set is scanned. Scan operation is supported only for Backup Exec server having Windows 2016 or later.

To modify the settings of the malware scan job, go to **Backup Exec Settings > Malware Scan** tab.

Note: Backup Exec uses the selected scan engine to scan the backup set chain for malware. Ensure that you have enabled the selected scan engine with the updated antivirus definitions.

See [“Configuring malware scan for virtual machine backup sets before restore”](#) on page 238.

See [“Changing malware scan options”](#) on page 783.

- 6 Click **Next**.
- 7 Select one or more transport modes to transport the vmdk file to the host.
You must select at least one of the transport mode options. If you select more than one option, the method is determined by the priority and the resources that are available. Click **Move Up** or **Move Down** to arrange the transport modes in the priority that you want to use.

NBD-Do not encrypt the virtual disk data for over-the-network transfers

Use this option if you do not use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

NBDSSL-Encrypt virtual disk data for over-the-network transfers

Use this option if you use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

Hotadd-Use virtual disk files from the Backup Exec server on the virtual machine

Use this option if you want to use the hotadd feature for ESX/ESXi. The hotadd feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hotadd.

SAN-Use the SAN to move virtual disk data

If you select this option, the virtual machine must have SAN read/write access to the VMware datastore that hosts the restore target. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi server is not affected.

Note: The SAN transport mode is not recommended for restores of thin provisioned disks because performance may be slower than the NBD transport mode.

- 8 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Delete existing virtual machines or selected disks prior to restore	Deletes existing virtual machine if you select the entire virtual machine or all the disks of a virtual machine. The virtual machine may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server unless you select this option. Deletes the selected disks if you select disks from within a virtual machine backup. Instead of deleting the virtual machine, only the selected disks are restored and the data from the remaining disks of the virtual machine is retained.
Power on virtual machine after restore	Select this option if you want Backup Exec to turn on the restored virtual machine after the restore job completes.

- 9 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 10 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 11 Review the job summary, and then click **Finish**.

To redirect the restore of VMware virtual machines or vmdk files to a different host

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup	Do the following in the order listed: ■ In the Details pane at the bottom of the screen, select the virtual machine. ■ Click Restore, and then select Restore GRT-enabled data. ■ In the Restore Wizard, select Files, folders, or volumes, and then click Next.

To restore the entire virtual machine or virtual disks Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **VMware data**, and then click **Next**.

- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
- 3 Select **To a different vCenter or ESX server**.
- 4 Select **Scan backup sets for malware before restore** to scan the selected backup set for malware before restore is performed.

For an incremental backup, the previous backup sets in the recovery chain are also used. To restore backup sets even if the malware scan fails, select the **Continue to restore even if scan fails** check box.

Malware scan operation is CPU intensive and can take significant time if the virtual machine backup has a large amount of data. In this scenario the time that the restore job requires can increase and you can schedule the scan. Scan operation is only supported for GRT enabled sets on disk and deduplication devices.

For GRT restore, if you are running a malware scan on backup sets as part of the restore job or ad hoc scan, the entire virtual machine backup set is scanned. Scan operation is supported only for Backup Exec server having Windows 2016 or later.

To modify the settings of the malware scan job, go to **Backup Exec Settings > Malware Scan** tab.

Note: Backup Exec uses the selected scan engine to scan the backup set chain for malware. Ensure that you have enabled the selected scan engine with the updated antivirus definitions.

See [“Configuring malware scan for virtual machine backup sets before restore”](#) on page 238.

See [“Changing malware scan options”](#) on page 783.

- 5 Click **Next**.
- 6 Type the name of the vCenter server or ESX/ESXi host that you want to restore to, or click **Browse** to search for the server.

- 7 Select the correct logon account for the server that you want to restore the data to.
- 8 Click **Browse** next to **Virtual machine folder** to select the folder to which you want to restore.
- 9 Click **Browse** next to **Resource pool** to select the resource pool to which you want to restore.
- 10 If you want to create a new name for the virtual machine, type it in the **Virtual machine name** field. You might want to enter a new name if a virtual machine with the same name already exists on the server.
- 11 Select the network that the new virtual machine should use after the restore job completes.
- 12 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Use the original disk datastore selections if available on the selected host Check this check box to use the original datastore selections on the virtual server. If the original datastore selections do not exist, then the datastore selections from the backup data are used.

Restore virtual machine to the most recent hardware version that the destination environment supports Check this check box to restore the virtual machine using the most up-to-date VMware hardware version that is available in the destination environment instead of using the virtual machine's original hardware version.

If you do not select this check box, then the virtual machine's original hardware version is preserved when it is restored.

Restore virtual clients with thin provisioning Check this check box to restore the virtual machine with thin provisioning. Thin provisioning can help you more efficiently dedicate storage capacity in your VMware ESX Server environment.

- 13 Select one or more transport modes to transport the vmdk file to the host.

You must select at least one of the transport mode options. If you select more than one option, the method is determined by the priority and the resources that are available. Click **Move Up** or **Move Down** to arrange the transport modes in the priority that you want to use.

NBD-Do not encrypt the virtual disk data for over-the-network transfers

Use this option if you do not use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

NBDSSL-Encrypt virtual disk data for over-the-network transfers

Use this option if you use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

Hotadd-Use virtual disk files from the Backup Exec server on the virtual machine

Use this option if you want to use the hotadd feature for ESX/ESXi. The hotadd feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hotadd.

SAN-Use the SAN to move virtual disk data

If you select this option, the virtual machine must have SAN read/write access to the VMware datastore that hosts the restore target. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi server is not affected.

Note: The SAN transport mode is not recommended for restores of thin provisioned disks because performance may be slower than the NDB transport mode.

- 14 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Delete existing virtual machines prior to restore Select this option if the virtual machine that you are restoring already exists on the server. If you select this option, the virtual machines may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server unless you select this option.

Power on virtual machine after restore Select this option if you want Backup Exec to turn on the restored virtual machine after the restore job completes.

- 15 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 16 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 17 Review the job summary, and then click **Finish**.

To restore to a different path

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled data**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **VMware data**, and then click **Next**.

- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
- 3 Select **To a different path**.
- 4 Enter the drive and path that you want to restore to.
- 5 Select **Scan backup sets for malware before restore** to scan the selected backup set for malware before restore is performed..

For an incremental backup, the previous backup sets in the recovery chain are also used. To restore backup sets even if the malware scan fails, select the **Continue to restore even if scan fails** check box.

Malware scan operation is CPU intensive and can take significant time if the virtual machine backup has a large amount of data. In this scenario the time that the restore job requires can increase and you can schedule the scan. Scan operation is only supported for GRT enabled sets on disk and deduplication devices.

For GRT restore, if you are running a malware scan on backup sets as part of the restore job or ad hoc scan, the entire virtual machine backup set is scanned. Scan operation is supported only for Backup Exec server having Windows 2016 or later.

To modify the settings of the malware scan job, go to **Backup Exec Settings > Malware Scan** tab.

Note: Backup Exec uses the selected scan engine to scan the backup set chain for malware. Ensure that you have enabled the selected scan engine with the updated antivirus definitions.

See ["Configuring malware scan for virtual machine backup sets before restore"](#) on page 238.

See ["Changing malware scan options"](#) on page 783.

- 6 Click **Next**.

- 7 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 8 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 9 Review the job summary, and then click **Finish**.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

About instant recovery of a VMware virtual machine

Backup Exec lets you recover a virtual machine instantly without waiting to transfer the virtual machine's data from a backup set. Backup Exec starts the instantly recovered virtual machine directly from the backup set and users can access it on the vCenter or ESX/ESXi host immediately. The startup time on the Backup Exec server depends on the network speed and storage speed, not on the size of the virtual machine.

To fully restore the virtual machine, use VMware Storage vMotion to migrate the virtual machine data files from the backup set to the vCenter or ESX/ESXi host. After you migrate the instantly recovered virtual machine, you can use the Agent for VMware to back up the virtual machine.

You can use an instantly recovered virtual machine to perform the same operations as a virtual machine. An instantly recovered virtual machine can be used to do the following:

- Access and restore individual files and folders on a virtual machine.
 - Test a patch on an instantly recovered virtual machine before you apply the patch to production systems.
 - Troubleshoot a virtual machine or host, such as when the production ESX host is unresponsive. You can use the instantly recovered virtual machine until the production system is back online.
 - Verify the backup set for a virtual machine.
 - Copy a vmdk file, and then remove the virtual machine.
 - Verify an application on a virtual machine.
 - Recover the virtual machine permanently with Storage vMotion.
- If you need to perform a disaster recovery, you can instantly recover a virtual machine and then schedule a migration to move it to permanent storage on

the vCenter or ESX/ESXi host. The instantly recovered virtual machine remains available during the migration, which decreases the amount of downtime.

Instantly recovered virtual machines use Backup Exec server storage. If you remove an instantly recovered virtual machine, any changes that you made to the virtual machine are lost. In order to avoid losing your changes, migrate the virtual machine from Backup Exec server storage if any changes are made to the instantly recovered virtual machine.

Because Backup Exec has recently enhanced the resiliency of Instant Recovery, changes that you made to the virtual machine are no longer lost if you restart the Backup Exec server or encounter a network connectivity issue. However, the virtual machine cannot be used until the server restart is complete or network connectivity is restored.

Ensure that you remove the virtual machine from Backup Exec server storage when it is no longer needed or has been migrated.

The following table describes the instant recovery process for a virtual machine.

Table C-3 Instant recovery process for a virtual machine

Step	Description
Step 1	You run an instant recovery job from a backup of a VMware virtual machine.
Step 2	Backup Exec virtualizes the backup set.
Step 3	Backup Exec creates a share as an NFS datastore on the Backup Exec server. The datastore becomes accessible to the host where the virtual machine will be instantly recovered. The host uses the NFS datastore as read-only. Note: Since Windows Server for NFS is not secure, the share that Backup Exec creates will be available on the network.
Step 4	Backup Exec creates a virtual machine on the host and configures the virtual machine with write access to the datastore.
Step 5	Backup Exec creates a snapshot of the virtual machine.
Step 6	Backup Exec starts up the virtual machine automatically if you select the option to power on the virtual machine.
Step 7	You can use Storage vMotion to migrate the virtual machine from Backup Exec server storage if you want to save any changes that were made to the virtual machine.

Table C-3 Instant recovery process for a virtual machine (*continued*)

Step	Description
Step 8	You can run a job to do one of the following: ■ Remove an instantly recovered virtual machine that you no longer need to use. ■ Remove the instantly recovered virtual machine from Backup Exec server storage after you use Storage vMotion to migrate the virtual machine. Warning: You cannot upgrade Backup Exec until you remove all of the instantly recovered virtual machines.

Differences between an instantly recovered virtual machine and a restored virtual machine

The following table describes the differences between an instantly recovered virtual machine and a restored virtual machine.

Table C-4 Differences between an instantly recovered virtual machine and a restored virtual machine

Item	Instantly recovered virtual machine	Restored virtual machine
Data transfer	Does not transfer virtual machine data to the instantly recovered virtual machine.	Transfers all data from the backup set to the host.
Job time	Instant recovery takes less time and jobs are not dependent on the virtual machine size.	Restore time is dependant on the size of the virtual machine and the network and storage speed.
Read and write operations	Uses the backup set image for all of the read operations. Uses a snapshot on the VMware host for all write operations.	Data is moved to the VMware host.
Data storage	Uses Backup Exec server storage until you migrate the instantly recovered virtual machine	Uses VMware host storage.

Table C-4 Differences between an instantly recovered virtual machine and a restored virtual machine (*continued*)

Item	Instantly recovered virtual machine	Restored virtual machine
Server restart	Because the Instant Recovery resiliency has been enhanced, if the Backup Exec server is restarted, the instantly recovered virtual machine remains accessible. As soon as the restart is complete or network is reconnected, the instantly recovered virtual machine is brought back online automatically. In the case of network connectivity issues, the instantly recovered virtual machine is accessible after connectivity is restored.	The Backup Exec server and the VMware host can be restarted.

See [“Requirements for instant recovery of a VMware virtual machine”](#) on page 1047.

See [“Notes about instant recovery of a VMware virtual machine”](#) on page 1048.

See [“Creating an instant recovery job for a VMware virtual machine”](#) on page 1049.

See [“Removing an instantly recovered VMware virtual machine”](#) on page 1052.

Requirements for instant recovery of a VMware virtual machine

Review the following requirements before you create an instant recovery job for a virtual machine:

- Install the Windows Server for NFS role on the Backup Exec server. For more information, see your Microsoft Windows documentation.
- Enable Backup Exec's Granular Recovery Technology to recover individual items from Microsoft applications in the VMware backup job. If your virtual machine is not eligible for GRT, you can enable the **Enable Instant Recovery of all virtual operating system types, even those not eligible for GRT** option. Note that enabling this option will change the storage format of your next full backup to a format that is compatible with instant recovery.
See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 1028.
- Use the virtual-based backup method to back up the virtual machine.

- Use disk storage for the VMware backup job.
Cloud storage and disk cartridge devices such as RDX are not supported.
- Ensure that you have sufficient disk storage space on the host to save any changes made to the instantly recovered virtual machine.
- Ensure that the host uses the supported Vsphere version for the Backup Exec release.
- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List.
You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

Notes about instant recovery of a VMware virtual machine

Review the following notes before you create an instant recovery job for a virtual machine:

- Backup Exec Management Command Line Interface (BEMCLI) is not supported for instant recovery jobs.
- The following apply in a CAS environment:
 - When backup sets are targeted to deduplication storage devices, the **Instant Recovery** option may be grayed out on Managed Backup Exec Server. It is recommended to run the instant recovery job on the server that has run the backup and duplicate job.
 - The central administration server or the managed Backup Exec server can instantly recover a virtual machine that is running on the server itself. However, only the managed Backup Exec server can remove the instantly recovered virtual machine that was recovered on the server.
 - You cannot administer an instant recovery job from a central administration server after you delegate the job to a managed Backup Exec server.
- Remove or migrate the instantly recovered virtual machine from Backup Exec server storage before you upgrade or uninstall Backup Exec. Upgrades, patch updates, and uninstalls are blocked on the Backup Exec server and the host on which the Agent for VMware is installed if an instantly recovered virtual machine is running.
- If a large number of instantly recovered virtual machines are running at the same time, it affects the performance of the Backup Exec server. You should periodically review the instantly recovered virtual machines that are running in your environment. Remove a virtual machine that you no longer need or

migrate the virtual machine to the host, and then remove the virtual machine from the Backup Exec server.

- Migration uses network bandwidth. You should perform the migration during a time when the bandwidth requirement of other processes is low.
- When you migrate an instantly recovered virtual machine, you should use a path on the VMware host that is different from the original location where the virtual machine was recovered.
- An instantly recovered virtual machine cannot be included in a backup job until it is migrated with VMware Storage vMotion.

Best Practices for instant recovery of a VMware virtual machine

Review the following best practice before you create an instant recovery job for a virtual machine:

- To perform Instant Recovery for an ESXi guest, Backup Exec creates an NFS share that has the required VMDK files and configures ESXi to use the NFS share. The permissions set on the NFS share are configured to limit access to requests originating from an ESXi IP address. If a guest virtual machine is set to use Network Address Translation (NAT) on ESXi, it can access the NFS share on the Backup Exec server. As a security best practice, it is recommended to configure port groups with specific NICs assigned to them so that virtual machines connected to one port group are separated from the other port group over which the ESXi has access to the NFS share.

Creating an instant recovery job for a VMware virtual machine

You can create an instant recovery job for a VMware virtual machine, and then recover the virtual machine to the original location or to an alternate location.

If you want to move the instantly recovered virtual machine from Backup Exec server storage to VMware storage, you can migrate the virtual machine data files or the required disks from the instantly recovered virtual machine to VMware storage using Storage vMotion. During the migration, the virtual machine data files are transferred to VMware storage while the virtual machine is still running.

Note: For applications such as SharePoint and Exchange, recover all of the virtual machines that the applications need to use to function properly. For example, to create a Microsoft Exchange environment, recover the virtual machines that run the Exchange client and the Active Directory, and then establish a connection between these two virtual machines.

See [“Requirements for instant recovery of a VMware virtual machine”](#) on page 1047.

Complete the following steps to create an instant recovery job for a VMware virtual machine.

To create an instant recovery job for a VMware virtual machine

- 1 On the **Backup and Restore** tab, select the virtual machine that you want to instantly recover.
- 2 In the **Instant Recovery** group, click **Instantly Recover a VM**.
- 3 On the **Instantly Recover a VM** dialog box, in the **Job name** field, type a name for the instant recovery job.
- 4 In the **Backup set selection** group box, in the **Show backup sets from** field, select the beginning and ending dates for the backup sets that you want to include in the backup set selection.

By default, only the backup sets for jobs that ran in the past 30 days are displayed.

- 5 In the **Disk-based backup set** field, select the backup set that you want to use for the instant recovery job.

Only backup sets that meet the requirements for instant recovery are included in the list.

- 6 In the left pane, select **Destination**, and then set the following options for the job:

Item	Description
Virtual machine name	Indicates the instantly recovered virtual machine name.
vCenter server or ESX host name	Indicates the name of the vCenter server or ESX host.
Server logon account	Uses the default logon account that appears. You can select another logon account to use for the vCenter server or ESX host.
Host	Indicates the name of the ESX host that will run the instantly recovered virtual machine.
Virtual machine folder	Indicates the name of the existing vSphere folder to which you want to instantly recover the virtual machine. The folder default is the root of the datacenter.
Resource pool	Indicates the name of the resource pool to which you want to instantly recover the virtual machine. The resource pool is optional.
VM network	Indicates the name of the network that the instantly recovered virtual machine should use after the job completes. Ensure that you select a network that is isolated from the source virtual machine. Otherwise, it may cause network conflicts between the source virtual machine and the instantly recovered virtual machine.
Datastore or datastore cluster for storing virtual machine writes	Indicates the name of the datastore where you want to store any temporary changes that are made to the instantly recovered virtual machine.
Power on the virtual machine after it is recovered	Enables Backup Exec to start the virtual machine automatically after it is instantly recovered.

- 7 On the **Instantly Recover a VM** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.

- 8 (Optional) On the **Instantly Recover a VM** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the job completes.
- 9 Click **OK**.

Removing an instantly recovered VMware virtual machine

You should remove an instantly recovered virtual machine from the Backup Exec server storage after it has been migrated or when you no longer need to use it. When you remove an instantly recovered virtual machine, Backup Exec removes all of the database entries and all of the folders that are created at the time of the instant recovery job.

If you remove the virtual machine from the VMware host using the VMware vSphere Client, you should also ensure that you remove the instantly recovered virtual machine from Backup Exec. When you remove an instantly recovered virtual machine from the VMware host, the configuration folders are deleted from the VMware storage. The NFS share that was created for the instantly recovered virtual machine on the Backup Exec server is also removed.

Note: Data lifecycle management (DLM) is postponed for the backup set that was used to create the instantly recovered virtual machine until you remove the virtual machine. The next cycle of DLM expires the backup set.

To remove an instantly recovered VMware virtual machine

- 1 On the **Backup and Restore** tab, select a vSphere server or an ESX host that contains the instantly recovered machine that you want to remove.

You can double-click a vSphere server or an ESX host, and then in the left pane, click **Recovered VMs** to view the virtual machines that were instantly recovered to the server.
- 2 In the **Instant Recovery** group, click **Remove a Recovered VM**, and then do one of the following:

To remove a recovered virtual machine now

Do the following in the order listed:

- Select **Use defaults and remove now**.
- On the **Select Recovered VMs to Remove** dialog box, select one or more recovered virtual machines to remove.
- Click **OK**.

To customize settings and then remove a recovered virtual machine Do the following in the order listed:

- Select **Customize settings and remove**.
- On the **Select Recovered VMs to Remove** dialog box, select one or more recovered virtual machines to schedule for removal.
- Click **OK**.
- On the **Remove a Recovered Virtual Machine** dialog box, in the **Job name** field, type a name for the job.
- In the **Server logon account** field, add or edit a logon account for the vCenter server or ESX host.
- Check the **Remove the virtual machine even though disks are on both Backup Exec server storage and the virtual machine host** check box to remove an instantly recovered virtual machine that is not fully migrated or has disks that are not on Backup Exec storage.
- On the **Remove a Recovered Virtual Machine** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.
- (Optional) On the **Remove a Recovered Virtual Machine** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the job completes.
- Click **OK**.

Troubleshooting the Agent for VMware

If you have trouble with VMware backup jobs, review the following issues and solutions.

Issue	Solution
My VMware backup job stalled without transferring any data.	Attempting to cancel a job in this state causes the job to go into a Cancel pending state indefinitely. Stop the beremote.exe process, and then restart all of the Backup Exec services.
In a vCenter 5.0 environment, backing up cloned virtual machines caused the Backup Exec Agent for Windows service to stop.	Edit and save the virtual machine configuration of the cloned virtual machine in the VMware vSphere Client before creating the backup job. You need to do this only one time on each cloned virtual machine.

About Recovery Ready for VMware virtual machines

With the **Recovery Ready** feature, you can use the **Validate VM for Recovery** operation to validate the recoverability of virtual machines. When you create and run a Validate Virtual Machine for Recovery job, it runs tests on the virtual machine after which it is marked as recovery ready.

Recovery ready virtual machines can be used to do the following:

- **Disaster recovery readiness:** In a disaster recovery scenario, administrators are sure that the validated virtual machines are recoverable.
- **Vault the backups to cloud or tape:** Validates the backup sets before you vault them to devices such as cloud and tape.
- **Audit and compliance of backups:** For meeting company audit and regulatory compliance requirements, you can provide validation information of your backups for virtual machines.

When the Validate Virtual Machine for Recovery job runs, the virtual machine is registered on the vCenter server or ESX host with the host name in the `Validate_VM_name_GUID` format and then powered on. The virtual machine is created only for validation purpose. It is a temporary virtual machine, which is removed after validation is completed. After the power on is complete, a heartbeat check is run to ensure that the VMware Tools Service is running.

During virtual machine validation there is no data transfer.

When you create the Validate Virtual Machine for recovery job, before the validation job runs, you can select the maximum time that is allowed for a virtual machine to boot. The default value is 10 minutes. You can select a value between 1 to 60 minutes.

See [“Configuring default backup settings”](#) on page 778.

After these checks are successfully completed, the virtual machine is marked as ready for recovery.

All information about the validation is part of the job log. After validation you can generate the **Recovery Ready Validation Summary** report to view the summary of the virtual machines that you validated.

See [“Recovery Ready Validation Summary”](#) on page 832.

To see the validation status, on the **Backup and Restore** tab, double-click the host name or on the **Storage** tab, double-click the disk storage name. In the left pane, click **Backup Sets**. For the selected server or storage view, the **Validation Status** is displayed.

The following table describes the process of validating VMware virtual machines for recovery.

Table C-5 Virtual machine validation process

Step	Description
Step 1	Add the vCenter server or ESX host.
Step 2	Take a disk-based GRT backup of the virtual machine that is hosted on the vCenter server or ESX host. The Recovery Ready feature supports full, incremental, and differential backup sets.
Step 3	Run a Validate Virtual Machine for Recovery job for a VMware virtual machine.
Step 4	Backup Exec virtualizes the backup set.
Step 5	Backup Exec creates a share as an NFS datastore on the Media server. The datastore becomes accessible on the host where the virtual machine is validated for recovery. The host uses the NFS datastore as read-only. Note: Since Windows Server for NFS is not secure, the share that is created by Backup Exec is available on the network.
Step 6	Backup Exec runs the following tests in the given order on the VMware virtual machine being validated. ■ Register the virtual machine ■ Power On the virtual machine ■ Heartbeat check of the virtual machine

Table C-5 Virtual machine validation process (*continued*)

Step	Description
Step 7	If all the tests are successful, Backup Exec marks the VMware virtual machine as validated.

Tests run on a VMware virtual machine being validated

The following tests are run on a VMware virtual machine that is being validated.

Table C-6 Tests run on a VMware virtual machine

Test	Description
Register the virtual machine	The first test is to register the virtual machine on the vCenter server or ESX host.
Power ON the virtual machine	The second test is to power ON the registered virtual machine.
Heartbeat test	The third test is to check the VMware Tools service. If this service runs successfully on the powered on virtual machine, the heartbeat test passes. VMware Tools must be installed and the VMware Tools service must be running on the backed up virtual machine for the heartbeat check to be successful. If VMware Tools is not installed, the validation job fails and for the corresponding backup set, Validation failed status is displayed.

All information related to the test results is also available in the job log.

Validation status for a VMware virtual machine

During validation, the following statuses are displayed for the backup set of a VMware virtual machine:

- **Validation successful:** The backup set has passed all tests and is recovery ready.
- **Validation failed:** The backup set has failed the tests and is not recovery ready.
- **Unable to validate:** The virtual machine is not validated because of environmental issues. Hence, the backup sets cannot be validated.

During validation the following statuses are displayed for the validation job of a VMware virtual machine:

- **Successful:** The validation job has passed all tests and the virtual machine is recovery ready.
- **Success with Exceptions:** The validation job has passed all tests but when there is a clean-up of resources that are used for validation, the power off fails.
- **Failed:** The validation job has failed the tests or validation is not performed.

If validation fails or Backup Exec is unable to validate a virtual machine for recovery, you can check the validation job log for more details.

The following matrix gives details about the tests and validation status.

Table C-7 Validation status for a VMware virtual machine

Validation test	Backup set status	Validation job status	Reason for failure
Register virtual machine	If registration is successful, go to the next check.		
	Validation failed	Failed	Backup set issue
	Unable to validate	Failed	Environment issue
Power ON virtual machine	If power ON is successful, go to the next check.		
	Unable to validate	Failed	Connectivity issue or check not attempted as registration failed
Heartbeat check	Unable to validate	Failed	Test not attempted
	Validation failed	Failed	Boot issue or VMware Tools issue
	Validation successful	Successful	Virtual machine is validated and recovery ready

See [“Requirements for validating a VMware virtual machine for recovery”](#) on page 1057.

See [“Notes about validating a VMware virtual machine for recovery”](#) on page 1058.

See [“Creating a validate virtual machine for recovery job”](#) on page 1059.

Requirements for validating a VMware virtual machine for recovery

Review the following requirements before you configure a Validate virtual machine for Recovery job for a VMware virtual machine:

- VMware Tools must be installed on the vCenter server or ESX host.
- Install the Windows Server for NFS role on the Media server. For more information, see your Microsoft Windows documentation.
- Enable Backup Exec's Granular Recovery Technology to recover individual items in the VMware backup job. Use the virtual-based backup method to back up the virtual machine. If your virtual machine is not eligible for GRT, you can enable the **Enable Instant Recovery of all virtual operating system types, even those not eligible for GRT** option. Note that enabling this option will change the storage format of your next full backup to a format that is compatible with instant recovery.
See ["Using Granular Recovery Technology \(GRT\) with the Agent for VMware"](#) on page 1028.
- Use disk storage for the VMware backup job. Cloud storage and disk cartridge devices such as RDX are not supported.
- Ensure that the host uses the supported Vsphere version for the Backup Exec release.
- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List. You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

Notes about validating a VMware virtual machine for recovery

Review the following notes before you create a Validate Virtual Machine for Recovery job for a VMware virtual machine:

- Backup Exec Management Command Line Interface (BEMCLI) is not supported for validate virtual machine for recovery jobs.
- The following applies in a CAS environment:
 - When backup sets are targeted to deduplication storage devices, the **Recovery Ready** option may be grayed out on Managed Backup Exec Server. It is recommended to run the recovery ready job on the server that has run the backup and duplicate job.
 - If the Central Administration Server (CAS) has a backup set, create the Validate Virtual Machine for recovery job only on the CAS server. If the Managed Backup Exec Server (MBES) has a backup set, create the Validate Virtual Machine for Recovery job only on the MBES server.

- You cannot administer a Validate Virtual Machine for Recovery job from a central administration server after you delegate the job to a managed Backup Exec server.

Best Practices for validating VMware virtual machine for recovery

Review the following best practices before you create a Validate Virtual Machine for Recovery job for a VMware virtual machine:

- While validating a virtual machine for recovery, VMware must be targeted to the same ESX host version or later.
- It is recommended that you have the latest VMware ISO tools for the backed up VMware virtual machine.
- The maximum number of Validate Virtual Machine for Recovery jobs that run in parallel depend on the available NFS mounts on the target ESX host. NFS mounts are also used by instantly recovered virtual machines. Ensure that there are required number of NFS mounts available to run multiple Validate Virtual Machine for Recovery jobs.
- While validating a virtual machine for recovery, Backup Exec creates an NFS share that has the required VMDK files and configures ESXi to use the NFS share. The permissions set on the NFS share are configured to limit access to requests originating from an ESXi IP address. If a guest virtual machine is set to use NAT on ESXi, it can access the NFS share on the Backup Exec server. As a security best practice, it is recommended to configure port groups with specific NICs assigned to them so that virtual machines connected to one port group are separated from the other port group over which the ESXi has access to the NFS share.

Creating a validate virtual machine for recovery job

You can create a Validate Virtual Machine for Recovery job for a VMware virtual machine.

See [“Requirements for validating a VMware virtual machine for recovery”](#) on page 1057.

Complete the following steps to create Validate Virtual Machine for Recovery job for a VMware virtual machine.

To create a validate virtual machine recovery job

- 1 On the **Backup and Restore** tab, select the virtual machine that you want to validate.
- 2 In the **Recovery Ready** group, click **Validate VM for Recovery**.

- 3
- On the **Validate Virtual Machine for Recovery** dialog box, in the **Job name** field, type a name for the validation job or use the default name.
- 4
- In the **Backup set selection** group box, select the backup sets that you want to include in the backup set selection.

Item	Description
Use the latest available disk-based backup set when the job is run	Indicates that only the latest available disk-based backup set is used when the VMware virtual machine validation recovery job is run.
Select a disk-based backup set	Indicates that an available disk-based backup set is used for the VMware virtual machine validation recovery job.
	Show backup sets from
	Indicates the beginning and ending dates for the backup sets that you want to include in the backup set selection.
	By default, only the backup sets for jobs that ran in the past 30 days are displayed.
Disk-based backup set	Indicates the backup set that you want to use for the validation job.
	Only backup sets that meet the requirements for validation are included in the list.

- On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Destination**, and then set the following options for the job.

Item	Description
vCenter server or ESX host name	Indicates the name of the vCenter server or ESX host. Note: You can also enter a vCenter server or host that is different from the source vCenter server or ESX host.
Server logon account	Uses the default logon account that appears. You can select another logon account to use for the vCenter server or ESX host.
Host	Indicates the name of the ESX host that runs the validated virtual machine.
Virtual machine folder	Indicates the name of the existing vSphere folder to which you want to validate the virtual machine. The folder default is the root of the datacenter.
Resource pool	Indicates the name of the resource pool to which you want to validate the virtual machine. The resource pool is optional.
Datastore or datastore cluster for storing virtual machine writes	Browse the name of the datastore where you want to store any temporary changes that are made to the validated virtual machine.

- On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.

Item	Description
Recurrence	Select this option to create a recurring schedule for the job.
Run now with no recurring schedule	Runs the job immediately without a recurring schedule.
Run on	Schedules the job to run on a specific date and time.
Create without a schedule	Creates the job without scheduling it. When you use this option, the job does not run at the time of creation. The job remains unscheduled until you choose to run it.
Submit job on hold	Enables you to submit the job, but not run the job until you change the job's hold status.

- (Optional) On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the Validate Virtual Machine for Recovery job completes.
- Click **OK**.

The Validate Virtual Machine for Recovery job runs and after it is completed successfully, the virtual machine is marked as recovery ready.

To see the validation status, on the **Backup and Restore** tab, double-click the vCenter server or ESX host name or on the **Storage** tab, double-click the disk storage name. In the left pane, click **Backup Sets**. For the selected server or storage view, you can view the status in the **Validation Status** column.

Backup Exec Agent for Microsoft Hyper-V

This appendix includes the following topics:

- [About the Agent for Microsoft Hyper-V](#)
- [Requirements for using the Agent for Microsoft Hyper-V](#)
- [About installing the Agent for Microsoft Hyper-V](#)
- [Notes about using the Agent for Hyper-V](#)
- [Disk space optimization with the Agent for Hyper-V](#)
- [Adding a Hyper-V host to the list of servers on the Backup and Restore tab](#)
- [Viewing details about Hyper-V resources](#)
- [Installing the Agent for Windows on Hyper-V virtual machines](#)
- [Push-installing the Agent for Windows to Hyper-V virtual machines](#)
- [Backing up Microsoft Hyper-V virtual machines](#)
- [Restoring Microsoft Hyper-V virtual machines](#)
- [About instant recovery of a Hyper-V virtual machine](#)
- [About Recovery Ready for Hyper-V virtual machines](#)
- [Troubleshooting issues with Backup Exec Agent for Microsoft Hyper-V](#)

About the Agent for Microsoft Hyper-V

The Backup Exec Agent for Microsoft Hyper-V (Agent for Hyper-V) lets you back up and restore the following resources:

- Hyper-V hosts that run on supported versions of Microsoft Windows.
You can find a list of supported operating systems, platforms, and applications in the Backup Exec Software Compatibility List.
- All virtual machines that reside locally on the Hyper-V hosts.
- Clustered Hyper-V hosts, including the virtual machines that reside on cluster shared volumes (CSV).
- The virtual machines that have their disks hosted on Server message block (SMB).
- The virtual machines that have their disks hosted on Scale-Out File Server.

The Hyper-V agent allows backup using three processing methods:

- The Resilient Change Tracking (RCT) method: This method is available for the virtual machines that are hosted on Microsoft Hyper-V Server 2016 or later and have configuration version of 8.0 or later. RCT is a native Microsoft Hyper-V mechanism for tracking changed blocks in the virtual hard disks of a virtual machine. During backup, Backup Exec requests Microsoft Hyper-V to create a checkpoint for the selected virtual machine. The backup is done for this checkpoint.

After the virtual machine backup, before the checkpoint is merged with the base virtual machine disk, Hyper-V converts the checkpoint to a reference point. The reference point represents the point-in-time view of the virtual machine disk state.

During the incremental backup of the same virtual machine, Backup Exec creates a new checkpoint and uses it as a source for the backup. Backup Exec queries Hyper-V to get the incremental changes between the reference point created during the previous run and the checkpoint created during the current run of the backup job. Only the changed data blocks are backed up from the created checkpoint.

Backup Exec supports both incremental backup and differential backup using the RCT method.

- The standard processing method: This method reads the whole virtual disk and identifies the changes that should be backed up. The changed blocks that are identified are then backed up. Backup Exec supports both incremental backup and differential backup using this method.
- The faster processing method: This method is faster than the standard processing method because it writes all changes to a new differencing disk

that is backed up. Backup Exec supports only incremental backup using this method. To configure Hyper-V incremental and differential backup settings please refer to the 'Configuring default backup settings' section.

See [“Configuring default backup settings”](#) on page 778.

For a new install

For a new installation of Backup Exec, the RCT method and the Standard processing method are selected by default.

For an upgrade install

For an upgrade install, the settings remain unchanged and the RCT method is not selected. The existing method that you select (standard or faster processing method) is not changed. When Backup Exec runs a backup for a virtual machine using the Hyper-V agent, following are the considerations to decide the method to use.

The Hyper-V version is Microsoft Windows Server 2016 or later and the selected virtual machine has configuration version 8.0 or later: Backup Exec attempts to use the RCT method if the RCT method is selected in Backup Exec Settings. If the RCT method is not selected, Backup Exec uses one of the selected methods (standard processing or faster processing).

If the Hyper-V version is Windows Server 2012 R2 or earlier, the RCT method is not supported. Backup Exec uses one of the selected methods between standard or faster processing.

Online and offline backups

Backup Exec can back up the virtual machines that are online or that are in an offline state or a saved state. The virtual machines that use Hyper-V Integration Services can be backed up while they are online. You can include both online and offline virtual machines in the same backup job. During the backup of an online virtual machine, Backup Exec takes a snapshot backup of the Hyper-V host. The host in turn takes a snapshot of the virtual machines on the host. This process enables Backup Exec to back up virtual servers without any downtime. If an online backup cannot be performed and the option **Exclude virtual machines that must be put in a saved state to back up** is selected for the backup job, then an offline backup is performed. With an offline backup, the virtual machine is placed briefly in a saved state. However, the virtual machine does not remain in the saved state for the entire backup job.

The amount of downtime for a saved state backup job depends on the following:

- The amount of memory that is allocated to the virtual machine.
- The current load on the host's operating system.

For information about the best practices to use Backup Exec Agent for Microsoft Hyper-V, refer to *Backup Exec Best Practices*.

See [“Requirements for using the Agent for Microsoft Hyper-V”](#) on page 1066.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 1074.

See [“Restoring Microsoft Hyper-V virtual machines”](#) on page 1093.

Requirements for using the Agent for Microsoft Hyper-V

You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

The following items are required:

Table D-1 Requirements for the Agent for Microsoft Hyper-V

Software	Installed on
A supported version of Microsoft Hyper-V	Microsoft Hyper-V host
Backup Exec	Backup Exec server Note: It is recommended that the Backup Exec server should run a version of Windows that is equal to or greater than the highest version of Windows that the virtual machines in your environment run. For example, if your virtual machines run Microsoft Windows Server 2012, the Backup Exec server should also run Microsoft Windows Server 2012.
Agent for VMware and Hyper-V	Microsoft Hyper-V host Note: If you use Hyper-V in a cluster environment, you must install the Agent for VMware and Hyper-V on each node in the cluster. If you do not, you may not be able to see all of the clustered backup sources that are available for backup.

Requirements for online backups

To run an online backup, the following requirements must be met:

- Check the Microsoft website for a Hyper-V Server version and its supported Windows server guest operating systems.
- Hyper-V Integration Services with Backup (Volume snapshot) is installed.

Note: Using an incorrect version of Hyper-V Integration Services may lead to a virtual machine being backed up offline instead of online. For example, if you copy a virtual machine from a Windows 2012 Hyper-V host to a Windows 2012 R2 Hyper-V host, the Hyper-V Integration Services are not updated automatically.

- The virtual machine is in a running state.

If these conditions are not met, the virtual machine is placed in a saved state if it is running. If the virtual machine is turned off, then that virtual machine is backed up only if you select the option **Back up virtual machines that are powered off** on the **Virtual Machines** options dialog box.

Requirements for collecting catalog data for Microsoft applications

To enable Backup Exec to collect catalog data for Microsoft Exchange, SharePoint, Active Directory, and SQL on the virtual machine, the following items are required on the virtual machine:

- A licensed version of the Backup Exec agent for the application.
- The Agent for Windows.
The Agent for Microsoft Hyper-V includes a license for the Agent for Windows. The agents for Microsoft Exchange, Active Directory, and SQL also include a license for the Agent for Windows. No separate license is required for the Agent for Windows.
- The virtual machine must be capable of being backed up online.
- The credentials that you use to access the virtual machine must also have access to the application.

Requirements for virtual machines

The following items are required for virtual machines:

- Virtual machines must have unique names. Backup Exec does not support virtual machines that have duplicate names.
- The Agent for Windows must be installed on the virtual machine to enable individual files and folders to be restored back to the original virtual machine.

- The Agent for Windows must be installed on the virtual machine to enable individual SQL databases, Exchange items, SharePoint items, and Active Directory objects to be restored back to the original virtual machine.
- Do not use special characters, such as the percent sign (%) or a forward slash (/) in Hyper-V virtual machine names. Special characters may cause redirected restores to fail.

See [“About the Agent for Microsoft Hyper-V”](#) on page 1064.

About installing the Agent for Microsoft Hyper-V

The Backup Exec Agent for Microsoft Hyper-V is installed as part of the Agent for VMware and Hyper-V. The Agent for Microsoft Hyper-V is installed on the Microsoft Hyper-V host. If your Backup Exec server is also your Microsoft Hyper-V host, you can install the Agent for Microsoft Hyper-V when you install Backup Exec. Or, you can install it after Backup Exec has been installed.

If Backup Exec is not installed on your Microsoft Hyper-V host, you must push-install the Agent for Windows to your Microsoft Hyper-V host. You do not need to install the Agent for Microsoft Hyper-V on virtual machines. However, a license is required on the Backup Exec server for the Agent for Microsoft Hyper-V. The Agent for Windows is included with the Agent for Microsoft Hyper-V.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 68.

Notes about using the Agent for Hyper-V

The Backup Exec Agent for Hyper-V lets you back up and restore Hyper-V virtual machines that are hosted on Microsoft Windows Server 2012 and later. This version of Backup Exec supports the Windows Server 2012 and later vhdx file format and Microsoft incremental backups.

- Virtual machines that have only vhd files
Both file/folder-level GRT and application-level GRT are fully supported.
- Virtual machines that have only vhdx files
 - For a Backup Exec server that runs Windows 2012 or later
Both file/folder-level GRT and application-level GRT are fully supported.
- Virtual machines that have a mixture of vhd and vhdx files

- For a Backup Exec server that runs Windows 2012 or later
Both file/folder-level GRT and application-level GRT are fully supported.
- Any virtual machines that are configured with Fiber Channel HBAs are skipped during backup jobs. The job log lists a message to indicate that the virtual machine was skipped.
- Remote VSS is not supported. In other words, virtual machines with vhd or vhdx files that are located on shares are not supported.
- A VSS Copy backup of a virtual machine is not supported.
- Backup Exec does not support instant recovery of a virtual machine to a Hyper-V server that runs an older version than the Hyper-V server from which you backed up the virtual machine.
- Any virtual machines that have storage spaces are not supported because Microsoft Hyper-V VSS is unable to take snapshots of virtual machines with storage spaces.
- Hyper-V Replication virtual machine backups may result in redundant backups of the primary virtual machine and the replicated virtual machine.
- Virtual machines that are configured with remote storage are skipped during backup jobs. The job log contains a message about the virtual machines that were skipped.
- If you back up a volume that is located on a vhdx disk and that was formatted with the "Perform a quick format" option deselected, then the size of the backup equals the size of the disk instead of the size of the data that was backed up.
- Virtual machines that have shared vhdx files are excluded from a backup job and the job fails.
- The restore of a Gen2 virtual machine can be redirected only to another Hyper-V host that runs Windows Server 2012 R2 or higher.
- To enable GRT for a Hyper-V virtual machine that runs Active Directory or Active Directory Lightweight and backup that virtual machine to tape, both the Hyper-V host and the Backup Exec server must run the same version of the Windows server. For example, if you want to enable GRT for a Hyper-V virtual machine that runs Active Directory or Active Directory Lightweight on Windows Server 2012 R2 and backup that virtual machine to tape, both the Hyper-V host and the Backup Exec server must run Windows Server 2012 R2.
- Backup Exec does not restore user defined checkpoints when a virtual machine is protected using the RCT method and the virtual machine had user defined checkpoints at the time of the backup. All of the data available in the virtual

machine disks including the checkpoint disks are restored but the checkpoints are not available in the Hyper-V manager.

If a virtual machine has user-defined checkpoints, Backup Exec does not restore the checkpoints when a virtual machine is protected using the RCT method.

- Restore of a virtual machine hosted on a SMB share/Scale Out File Server share can only be done to the original location or to a locally hosted volume on the Hyper-V host. Redirected restore to a different SMB share/Scale Out File Server share is not supported. After a virtual machine is restored, It is recommended that you run a full backup of that virtual machine.

See [“Requirements for using the Agent for Microsoft Hyper-V”](#) on page 1066.

Disk space optimization with the Agent for Hyper-V

Backup Exec attempts to perform disk space optimization during Hyper-V backup jobs. Disk space optimization is performed whenever possible.

However, disk space optimization cannot be performed in the following situations:

- The file system is not NTFS. For example, if the file system is FAT, ReFS, or Linux, then disk space optimization is not performed.
- The disks are configured with Windows Storage Spaces.
- The volumes are configured with dynamic volumes, such as spanned, striped, mirrored, or RAID 5, for either MBR or GPT disks.
- The vhd file or vhdx file is not consistent at the time of backup.

A file may not be consistent for the following reasons:

- If an offline backup of a running virtual machine is performed. This may occur if the Hyper-V Integration Services are not installed or if the version of Integration Services running on a virtual machine do not match the version running on a Hyper-V host. This may also occur if the virtual machine runs an operating system that is not supported.
- The user-created checkpoints (snapshots) do not generate consistent disks.

Adding a Hyper-V host to the list of servers on the Backup and Restore tab

You can add a Hyper-V host to the list of servers on the **Backup and Restore** tab so that the host server and the virtual machines that it hosts can be selected for backup jobs.

To add a Hyper-V host to the list of servers on the Backup and Restore tab

- 1

On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2

Select **Microsoft Hyper-V host**, and then click **Next**.
- 3

Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4

Select **Add a Microsoft Hyper-V server to the list of servers**.
- 5

Type the name or IP address of the Hyper-V host that you want to add, and type an optional description.
- 6

If you want to install the Agent for Windows on all of the host's virtual machines, verify that the option **After adding the virtual host, install the Backup Exec Agent for Windows on the virtual machines** is selected.

If you do not want to install the Agent for Windows on the guest virtual machines, clear the check box.
- 7

Click **Next**.
- 8

Select the correct logon account for the server from the drop-down list.

If the correct logon account is not in the list, click **Add/Edit** to add it to the list.
- 9

Click **Next**.
- 10

Select any of the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically	Enables Backup Exec to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.
Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required	Enables Backup Exec to automatically restart the remote computer, if required.

- 11

Review the summary information, and then click **Install**.

Viewing details about Hyper-V resources

The **Backup and Restore** tab includes a **Details** pane on the lower portion of the screen. The **Details** pane provides additional details for any type of server that

is selected from the list of servers in the upper portion of the screen. Additional details and functionality appear if a Hyper-V host is selected. The **Details** pane for Hyper-V resources includes details about the last 7 days of backup jobs, the date of the last backup, and the date of the next scheduled backup. In addition, it includes the ability to back up and restore data, and to filter the list of guest virtual machines.

The **Details** pane for Hyper-V virtual machines includes the resources that may not appear in the list of servers, such as:

- Virtual machines that do not have the Agent for Windows installed.
- Non-Windows virtual machines, such as Linux.
- Virtual machines that do not have a GRT-enabled backup.

The **Details** pane for Hyper-V virtual machines provides a combination of current information and historical information. The **Refresh** button enables Backup Exec to discover all server resources. However, if a virtual machine has been moved, deleted, or has had a credentials change, then the following occurs:

- If the virtual machine has never been backed up, Backup Exec removes it from the **Details** pane.
- If the virtual machine has been backed up, Backup Exec does not remove it from the **Details** pane.

If a virtual machine is renamed, then the following occurs in the **Details** pane:

- If the virtual machine has been backed up, the **Details** pane includes an entry for the new name and keeps the entry for the old name.
- If the virtual machine has never been backed up, the **Details** pane includes only the new virtual machine name.

To view details about Hyper-V resources

- ◆ On the **Backup and Restore** tab, select a Hyper-V host from the list of servers.

The resources for the selected host appear in the **Details** pane.

Installing the Agent for Windows on Hyper-V virtual machines

To use Backup Exec's Granular Recovery Technology (GRT) for Microsoft application data, install the Agent for Windows on any virtual machines that run Windows.

To install the Agent for Windows on Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Hyper-V host**, and then click **Next**.
- 3 Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Select **Install the Backup Exec Agent for Windows on the virtual machines of a Microsoft Hyper-V host**.
- 5 Select the host server from the drop-down list, and then click **Next**.
- 6 Select the virtual machines that you want to install the Agent for Windows on, and then click **Next**.
- 7 Select the logon account for the virtual machines, and then click **Next**.
- 8 Select any of the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically

Enables Backup Exec to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required

Enables Backup Exec to automatically restart the remote computer, if required.

- 9 Review the summary, and then click **Install**.

Push-installing the Agent for Windows to Hyper-V virtual machines

You use the **Add Server** wizard to push-install the Agent for Windows to the Hyper-V virtual machines that you select. You can choose to install the Agent for Windows on all of the guest virtual machines that need it or on specific virtual machines only. For example, if you know that some of your guest virtual machines are scratch machines and do not need to be backed up, you can exclude those virtual machines from the installation.

To push-install the Agent for Windows to Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, right-click the virtual host, and then select **Install the Agent for Windows to guest virtual machines**.
- 2 Check the option **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the guest virtual machines of a Microsoft Hyper-V server**.
- 4 In the **Microsoft Hyper-V server** field, select the Hyper-V server that hosts the virtual machine.
- 5 Click **Next**.
- 6 Select the guest virtual machines on which you want to install the Agent for Windows, and then click **Next**.
- 7 Verify that the virtual machines you selected are online and select the appropriate logon account for those virtual machines, and then click **Next**.
- 8 Review the information on the **Summary** panel, and then click **Install**.

Backing up Microsoft Hyper-V virtual machines

The following backup selections are available for Microsoft Hyper-V:

Table D-2 Microsoft Hyper-V backup selections

Container name	Items in the container	What is included in the backup job
Microsoft Hyper-V	This item includes Initial Store and Virtual Machines .	If you select the Microsoft Hyper-V container for backup, the backup job includes the application configuration settings and all virtual machines.
Initial Store Note: This container does not appear for Hyper-V 2012 and later.	This item includes the virtual server application configuration settings.	If you select Initial Store for backup, the backup job includes a single XML file that contains the Hyper-V authorization configuration.
Host Component Note: This container appears only for Hyper-V 2012 and later.	This item includes the virtual server application configuration settings.	If you select Host Component for backup, the backup job includes multiple files that contain the Hyper-V authorization configuration.

Table D-2 Microsoft Hyper-V backup selections (continued)

Container name	Items in the container	What is included in the backup job
Virtual Machines	This item includes each virtual machine that resides on the virtual server. Note: The parent disks for the Hyper-V virtual machines are displayed. When you select a parent disk, Backup Exec backs up the selected parent and the respective child disks of the virtual machine.	

Table D-2 Microsoft Hyper-V backup selections (*continued*)

Container name	Items in the container	What is included in the backup job
		For an individual virtual machine, based on the disks that you select or if you select the entire virtual machine, the backup includes the following items: ■ vhd files for the selected disks ■ avhd files for the selected disks ■ Memory snapshot files ■ vsv files (not present in Hyper-V 2016) ■ bin files (not present in Hyper-V 2016) ■ vhdx files (Hyper-V 2012 and later) for the selected disks ■ avhdx files (Hyper-V 2012 and later) for the selected disks ■ .vmcx files (Hyper-V 2016) ■ The virtual machine's configuration *.XML files ■ Hyper-V managed snapshots Note: If you attempt to back up a virtual machine with the NetBIOS name "LocalHost", the backup will fail with the error "e000976f - Cannot backup the virtual machine to the deduplication device that is hosted by the same virtual machine." Note: Virtual machines that have remote vhd files are excluded from the backup job. You can use the Agent for Windows and the appropriate Backup Exec agent to protect virtual machines that have remote vhd files. Warning: Backup jobs fail for virtual machines that have pass thru disks. You can install the Agent for Windows and the appropriate Backup Exec agent on virtual machines that have pass thru disks and protect them as if

Table D-2 Microsoft Hyper-V backup selections *(continued)*

Container name	Items in the container	What is included in the backup job
		they are physical computers. Note: You may experience a processing time of more than 30 minutes if you back up a virtual machine that runs any supported version of Microsoft SharePoint.

Note: If you want to back up an Exchange database availability group (DAG) on a virtual machine, use the Backup Exec Agent for Exchange. The Agent for Hyper-V does not support backups of Exchange DAG.

To back up Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, select a virtual machine or a Hyper-V host from the list of servers.
- 2 Do one of the following:

If you selected a Hyper-V host in step 1

In the **Backups** group at the top of the screen, click **Backup**, and then select the type of backup you want to perform.

Alternatively, in the **Details** pane at the bottom of the screen, select the resources that you want to back up, and then click **Backup**. Select the type of backup you want to perform.

If you selected a virtual machine in step 1 Do the following:

- In the **Backups** group at the top of the screen, click **Backup**, and then select the type of backup that you want to perform.
 - On the **Back Up Virtual Machine** dialog box, select the backup method that you want to use for this virtual machine, either **Virtual-based backup** or **Agent-based backup**, and then click **Next**.
See [“Recommendations for when to use virtual-based backup and agent-based backup”](#) on page 139.
 - If you selected **Virtual-based backup**, select the virtual machine's host, and then click **Next**.
If the virtual machine's host is not in the list of servers, click **Add**, and then complete the steps in the wizard to add the host.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit** to add or remove resources from the backup selection list.
 - 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.
 - 5 Select the disks for a virtual machine that you want to back up.
You can select the entire virtual machine (select all disks) or select disks individually. You can also exclude disks selectively.

6 Click **OK**.

If you make partial selections, the **Virtual machines have partially selected disks** pop-up is displayed.

The job selections consist of one or more virtual machines where only some of its virtual disks are selected. For a virtual machine to function at the operating system level and application level, ensure that all the required disks are selected.

To perform File/Folder GRT and application GRT restore, ensure that the system disk is selected. For application GRT, select all virtual disks with application data.

See [“Configuring malware scan for virtual machine backup sets before restore”](#) on page 238.

See [“Changing malware scan options”](#) on page 783.

7 Click **OK**.

8 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.

9 On the **Backup Options** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.

10 On the **Backup Options** dialog box, in the left pane, select **Virtual Machines**.

11 Set any of the following options for this job.

Maximum number of streams used for parallel backup Select the maximum number of streams you want to use for parallel backup of virtual disks of a virtual machine.

You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage.

Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups.

See [“Multi-stream backup of virtual machines”](#) on page 1033.

Minimum number of streams used for parallel backup Select the minimum number of streams you want to use for parallel backup of virtual disks of a virtual machine.

You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage.

The concurrent write sessions value for the device must be greater than the minimum number of streams that are used for parallel backups, or the job fails.

Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups.

See [“Multi-stream backup of virtual machines”](#) on page 1033.

Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to run a full backup job if an incremental backup or a differential backup of the virtual machine cannot be performed. Backup Exec may not be able to perform an incremental backup or a differential backup for a number of reasons, such as if the snapshot configuration is altered or the configuration of the host server changed. If this option is not selected and an incremental backup or a differential backup cannot be performed, the job fails.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines when they are not powered on.
Enable Instant Recovery for all virtual machines, even those not eligible for GRT	Select this option to enable Instant Recovery of all virtual machines, even those that run operating systems which are not eligible for Granular Restore Technology.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the full backup. You must install the Agent for Windows on the virtual machine on which you want to restore the data. The Agent for Windows does not have to be installed on the virtual machine to back up the data.
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Exchange is installed.

Enable GRT for Microsoft SQL (database-level only) on virtual machines

Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SQL is installed.

Run an SQL Log backup after backing up the virtual machines

Select this option to enable Backup Exec to back up the SQL log. This option applies to SQL databases that update files in a log instead of in the database file. After the SQL log is backed up, the data from the log files is committed to the database and the log files are emptied. If you do not select this option, the log file continues to grow until it becomes full or until you manually back it up.

Enable GRT for Microsoft SharePoint on virtual machines

Select this option to enable Backup Exec to collect the information that is required to restore SharePoint data on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SharePoint is installed.

Exclude virtual machines that must be put in a saved state to back up

Select this option to exclude from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.

Backup method

Select this option to change the backup method for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties.

Note: Backup method is not applicable for Forever Incremental backup jobs.

- 12 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
- 13 Click **OK**.
- 14 On the **Backup Definition Properties** dialog box, click **OK**.

If you selected the virtual-based backup method, the backup job appears in the jobs list under the virtual host's name or IP address.

Setting default backup options for Hyper-V

You can use the defaults that Backup Exec sets during installation for all Hyper-V backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Hyper-V

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for virtual machine backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select **Virtual Machines**.

4 Select the default options that you want to use.

Maximum number of streams used for parallel backup

Select the maximum number of streams you want to use for parallel backup of virtual disks of a virtual machine.

You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage.

Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups.

See [“Multi-stream backup of virtual machines”](#) on page 1033.

Minimum number of streams used for parallel backup

Select the minimum number of streams you want to use for parallel backup of virtual disks of a virtual machine.

You can change the backup streams as required. Each stream performs the backup of a virtual disk independently. You can configure parallel streams only for backup-to-disk and deduplication disk storage.

The concurrent write sessions value for the device must be greater than the minimum number of streams that are used for parallel backups, or the job fails.

Multistream backups are not enabled by default for the backup jobs. The default value is 1, which means that single-stream backup is enabled. You must change the stream count to enable multistream backups.

See [“Multi-stream backup of virtual machines”](#) on page 1033.

Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to run a full backup job if an incremental backup or a differential backup of the virtual machine cannot be performed. Backup Exec may not be able to perform an incremental backup or a differential backup for a number of reasons, such as if the snapshot configuration is altered or the configuration of the host server changed. If this option is not selected and an incremental backup or a differential backup cannot be performed, the job fails.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines when they are not powered on.
Enable Instant Recovery for all virtual machines, even those not eligible for GRT	Select this option to enable Instant Recovery of all virtual machines, even those that run operating systems which are not eligible for Granular Restore Technology.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the full backup. You must install the Agent for Windows on the virtual machine on which you want to restore the data. The Agent for Windows does not have to be installed on the virtual machine to back up the data.
Enable GRT for Microsoft Active Directory objects on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Exchange is installed.

Enable GRT for Microsoft SQL (database-level only) on virtual machines

Enables Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SQL is installed.

Run an SQL Log backup after backing up the virtual machines

Enables Backup Exec to back up the SQL log. This option applies to SQL databases that update files in a log instead of in the database file. After the SQL log is backed up, the data from the log files is committed to the database and the log files are emptied. If you do not select this option, the log file continues to grow until it becomes full or until you manually back it up.

Enable GRT for Microsoft SharePoint on virtual machines

Enables Backup Exec to collect the information that is required to restore SharePoint data on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SharePoint is installed.

Exclude virtual machines that must be put in a saved state to back up

Excludes from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.

Backup method

Enables you to change the backup method for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties.

Note: Backup method is not applicable for Forever Incremental backup jobs.

5 Click **OK**.

How Backup Exec automatically protects new virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added

between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

In the backup selection list, dynamic inclusion is enabled for the following Hyper-V nodes:

- Microsoft Hyper-V
- Virtual Machines under Microsoft Hyper-V
- The Hyper-V host node
If you select the host node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V node.
- Microsoft Hyper-V HA Virtual Machines
- The cluster name node
If you select the cluster name node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V HA Virtual Machines node.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 1074.

Using Granular Recovery Technology (GRT) with the Agent for Hyper-V

Backup Exec Granular Recovery Technology (GRT) lets you restore individual files and folders without having to restore the entire virtual machine. It also lets you restore individual items from the following VSS-aware applications that are installed on virtual machines.

Backup Exec performs a single-pass backup to protect the host configuration data, all virtual machines, and VSS-aware applications that are installed on the virtual machines. Backup Exec's file/folder-level Granular Recovery Technology (GRT) is enabled by default for backup jobs. You can use a GRT-enabled backup to restore individual files and folders from a Windows virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange, SharePoint, and Active Directory applications that reside on virtual machines. You can also restore individual databases from Microsoft SQL when it resides on virtual machines.

Note: You must have the appropriate Backup Exec agent for Microsoft Exchange, SQL, SharePoint, or Active Directory on the virtual machine to perform GRT.

Table D-3 Types of data that Backup Exec protects for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec protects
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units
Microsoft SharePoint	SharePoint databases

Note: GRT is not intended to be used for system recovery. However, you can perform a complete system recovery by selecting the entire virtual machine as a restore selection in a restore job.

When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using GRT. By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

To use GRT, you must select the individual files and folders that you want to restore from the list that appears when you expand the Netbios name or the computer name of the virtual machine. You cannot select individual folders and files from the virtual machines that appear when you expand the **Virtual Machines** node.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

During the backup job, Backup Exec collects metadata for the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

Backup Exec cannot collect metadata in the following situations:

- GRT is disabled for an application.
- Backup Exec cannot connect to the virtual machine.

- Incorrect credentials were entered for the virtual machine.

Note: Backup Exec uses the Microsoft Hyper-V writer during backups of VSS-aware applications on virtual machines. The Microsoft Hyper-V writer truncates application logs before data is moved to the storage device. Therefore, the application logs for the applications on the virtual machines are truncated if you use Microsoft Hyper-V.

Requirements for using GRT to back up Microsoft application data on virtual machines

The following items are required to protect data for Microsoft Exchange, SQL, Active Directory, and SharePoint on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.
- The Backup Exec server must be able to connect to the virtual machine using the network name or IP address.
- The Backup Exec Agent for Windows must be installed on the virtual machine.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.
- The virtual machine cannot use dynamic disks, such as spanned, mirrored, striped, or RAID 5.

Unsupported configurations for GRT

Table D-4 Unsupported configurations for GRT

Unsupported items	Details
Restores of full and incremental backup sets from different storage devices	Backup Exec does not support restores from mixed media if GRT was enabled in the backup job. For example, if the full backup is on tape and the incremental backup is on a disk storage device, the restore job will fail. Restores from mixed media types are supported if GRT is not enabled.

Table D-4 Unsupported configurations for GRT (*continued*)

Unsupported items	Details
Virtual machines that have dynamic disks (with MBR or GPT partition style)	Backup Exec does not support granular recovery of files, folders, and applications on virtual machines that have dynamic disks (with GPT or MBR partition style).
Virtual machines that have ReFS and Deduplicated volumes	Backup Exec supports file/folder GRT for ReFS and Deduplicated volumes or application-level GRT for a virtual machine only if the Backup Exec server and HyperV host run the same or higher version of the operating system than the virtual machine. For Deduplicated volumes, the Backup Exec server must have the Windows Deduplication role/feature installed. For example, if the virtual machine runs on Windows 2016 and has REFS/Deduplicated volumes, then the Backup Exec server and HyperV host should also run on Windows 2016 or later. For Deduplicated volumes, Backup Exec server must have the Windows Deduplication role/feature installed.
VHDX format of virtual disks that are larger than 2 TB	File/folder-level and application level-GRT for a virtual machine that contains one or more VHDX files larger than 2 TB is not supported if the Backup Exec server is not running on Windows 2012 or later.

About backing up and restoring highly available Hyper-V virtual machines

When virtual machines are configured for high availability, they appear in the **Highly Available Hyper-V Machines** node in the backup selection list. Virtual machines that are not configured for high availability remain in the **Microsoft Hyper-V** node. When you make a backup selection, Backup Exec checks for highly available virtual machines. If highly available virtual machines are discovered, Backup Exec reminds you to select those virtual machines for backup.

The restore selections are similar to the backup selections. You can restore a highly available virtual machine in the same way you restore any other virtual machine. The virtual machine maintains its high availability. However, if you redirect the restore to another Hyper-V host, then the virtual machine is no

longer highly available when the restore job completes. You must reconfigure the virtual machine to be highly available.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 1074.

About backing up VMs hosted on SMB/Scale out File Server

Backup Exec supports backup of virtual machines hosted on SMB and Scale Out File Servers when the virtual machines are hosted on Microsoft Hyper-V Server 2016 or later and have configuration version of 8.0 or later. Backup Exec uses the Resilient Change Tracking (RCT) method to protect such virtual machines. The permissions for these virtual machines must be configured as per Microsoft documentation

The logon account specified in the backup job should have administrative access on the Hyper-V host and also have full permissions on the SMB share/Scale Out File Server share.

During the backup or restore of the virtual machine, Backup Exec processes the reads and writes from the Hyper-V host that owns the virtual machine. The data flow from the share goes first to the Hyper-V host and then to the Backup Exec server.

See [“About the Agent for Microsoft Hyper-V”](#) on page 1064.

How cataloging works with Hyper-V virtual machine backups

When you enable Granular Recovery Technology (GRT) for a backup job of a virtual machine, you can choose to run the catalog job for GRT as part of the backup job, as a separate job immediately after the backup job completes, or according to a schedule. By default, the catalog operation runs immediately after the backup job completes.

Note: The Instant GRT or full catalog features are not supported for backups to tape.

The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. You may want to schedule the catalog operation to run outside of your backup window so that it does not interfere with backup jobs. If the catalog operation is scheduled, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery on Hyper-V virtual machines. Before the full catalog job completes, instead of using the Search wizard, you must browse the backup sets to select the items that you want to restore.

For example, if you set up incremental backups to run every 11 hours and set up the catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)
- Incremental 1 (10:00 P.M.)
- Catalog 1 (Midnight). This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)
- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight). This job catalogs Incremental 5. Incremental 4 is not cataloged.
- Incremental 6 (5:00 A.M.). This backup is not cataloged.

In the example, the full catalog operation runs only for Incremental 5, Incremental 3, and Incremental 1. For such jobs, you can use the Search wizard to search the data or you can quickly browse for individual items that you want to restore. You can perform a granular recovery using Incremental 2, Incremental 4, and Incremental 6 as well; however, it takes slightly longer to browse items because they are not fully cataloged. Backup Exec dynamically displays the granular data by mounting the backup set.

How byte count is calculated for catalog operations

On the **Job Monitor** and **Job History**, the byte count that displays for a catalog operation may differ from the byte count that displays for the corresponding backup job. The byte count for a catalog job may be larger than the byte count for a backup job. The way in which Backup Exec catalogs the data affects the byte count that appears for the catalog job.

- When a catalog operation is performed for a full backup, the data is read on a file-by-file basis and the byte count is calculated accordingly. During the full backup job, the data is read in terms of the number of sectors and the byte count is calculated based on the number of sectors. Therefore, the byte count for the catalog job may be larger than the byte count for the backup job.
- When the catalog operation is performed for an incremental backup, all files on the virtual disk are cataloged instead of only the changed files. Therefore,

the byte count for the catalog job takes into account both the full backup and the incremental backup.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Restoring Microsoft Hyper-V virtual machines

You can use the Restore Wizard to restore data from virtual machines in the following ways:

- Restore a complete virtual machine for disaster recovery purposes.
- Restore individual files or folders that were backed up from the virtual machine (if you selected the Granular Recovery Technology features for the backup job).
- Restore a virtual machine to a different Microsoft Hyper-V server.
- Redirect flat files from the virtual machine to any computer that has an Agent for Windows installed.

Notes about restoring Hyper-V virtual machines

- Linux virtual machines must be restored in their entirety at the vhd level.
- The restore of a Microsoft Hyper-V virtual machine that is created within a mount point fails if the mount point does not exist at the time of the restore. To avoid this issue, configure the virtual machine to use volume GUID paths with no mount points.
- Restores of virtual machines with pass-thru disks, fibre channel adapters, and shared vhdx files may fail if those items have been deleted or those items are unavailable at the time of restore. The restore job will succeed if you have not deleted the pass-thru disks, fibre channel adapter, or shared vhdx file.
- Redirected restores of partially-selected virtual machines can be performed, but the virtual machine will not be registered. Backup Exec attempts to register virtual machines only if they are restored in full.
- When you restore a virtual machine that has excluded disks, Backup Exec removes the disks from the VM and deletes them from the host. In such a scenario, one of the following is possible:
 - Backup Exec successfully removed the disk from the virtual machine.
 - Backup Exec successfully removed the disk from the virtual machine but errors occurred when deleting the disk from the host. You can manually delete the disks from the host.

- Backup Exec cannot remove the disk from the virtual machine as some errors occurred. You can manually remove the disks from the virtual machine.

Note: The steps below apply to virtual machines that were backed up using the virtual-based backup method. If you backed up the virtual machine using the agent-based backup method, follow the steps for restoring a non-virtual backup

See [“Methods for restoring data in Backup Exec”](#) on page 235.

To restore Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, do one of the following.

To restore individual files and folders from a GRT-enabled backup

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled backup**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **Hyper-V data**, and then click **Next**.

- 2 Select the data that you want to restore, and then click **Next**.

Note: For virtual machines that have a System Reserved partition and are backed up with Granular Recovery Technology enabled, Backup Exec displays the virtual machine under a volume GUID instead of a drive letter. To restore data from these virtual machines, select one of the options to restore to another location.

3 Select the location where you want to restore the data.

To the original location	Select this option to restore the virtual machine to the same location from which it was backed up.
To a different server keeping the same drive and path location	Select this option to restore the virtual machine to a different server, but use the same drive and path name that were used in the original. You must enter the name of the server to which you want to redirect the restore, and you must also enter the logon account for that server Example: The original virtual machine was on \\ServerA\D:\VMs\1.vhd and you want to restore it to \\ServerB\D:\VMs\1.vhd
To a different server keeping the same path, while changing the drive	Select this option to restore the virtual machine to a different server and change the drive, but use the same path name that was used in the original. You must enter the name of the server and the drive to which you want to redirect the restore. You must also enter the logon account for that server. Example: The original virtual machine was on \\ServerA\D:\VMs\1.vhd and you want to restore it to \\ServerV\E:\VMs\1.vhd
To a different server: prefix the source drive and path with a new drive and path	Select this option to restore the virtual machine to a different server, change the drive, and include the original drive and path names in the new path. You must enter the name of the server, the drive, and the path to which you want to redirect the restore. You must also enter the logon account for that server. Example: The original virtual machines were on \\ServerA\D:\VMs\1.vhd and \\ServerA\E:\VMs\2.vhd, and you want to restore them to \\ServerB\Z:\ReplicatedVMs\D\VMs\1.vhd and \\ServerB\Z:\ReplicatedVMs\E\VMs\2.vhd

- 4 Select the register related check box to register the virtual machine on the Hyper-V server after it is restored to the new location.
- 5 Select **Scan backup sets for malware before restore** to scan the selected backup set for malware before restore is performed.

For an incremental backup, the previous backup sets in the recovery chain are also used. To restore backup sets even if the malware scan fails, select the **Continue to restore even if scan fails** check box.

Malware scan operation is CPU intensive and can take significant time if the virtual machine backup has a large amount of data. In this scenario the time that the restore job requires can increase and you can schedule the scan. Scan operation is only supported for GRT enabled sets on disk and deduplication devices.

For GRT restore, if you are running a malware scan on backup sets as part of the restore job or ad hoc scan, the entire virtual machine backup set is scanned. Scan operation is supported only for Backup Exec server having Windows 2016 or later.

To the modify settings of the malware scan job, go to **Backup Exec Settings > Malware Scan** tab.

Note: Backup Exec uses the selected scan engine to scan the backup set chain for malware. Ensure that you have enabled the selected scan engine with the updated antivirus definitions.

See [“Configuring malware scan for virtual machine backup sets before restore”](#) on page 238.

See [“Changing malware scan options”](#) on page 783.

- 6 Click **Next**.

- 7 Select the additional options that you want to use for this restore job, and then click **Next**.

Overwrite powered on virtual machines Select this option to enable powered on virtual machines to be overwritten and restored. By default, virtual machines must be turned off before the restore job processes and the virtual machine is overwritten. If a virtual machine is running during a restore job, but this option is not selected, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.

Do not overwrite powered on virtual machines Select this option to prevent powered on virtual machines from being overwritten and restored. By default, virtual machines are turned off before the restore job processes and the virtual machine is overwritten. If a virtual machine is running during a restore job, and this option is selected, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.

Power on the virtual machine after restore and resume from the available saved state Select this option to turn on the virtual machine automatically after the restore job completes. The virtual machine resumes operations from the saved state from the time of the backup.

Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

Power on the virtual machine after restore and discard the available saved state

Select this option to turn on the virtual machine automatically after the restore job completes. The virtual machine discards the available saved state.

Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

- 8 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 9 Review the job summary, and then click **Finish**.

About instant recovery of a Hyper-V virtual machine

Backup Exec lets you recover a virtual machine instantly from a backup set without waiting to transfer the virtual machine's data. Backup Exec starts the instantly recovered virtual machine directly from the backup set and users can access it on the Hyper-V host immediately. The startup time depends on the network speed and storage speed and not on the size of the virtual machine. You can use an instantly recovered virtual machine to perform the same operations as a virtual machine.

An instantly recovered virtual machine can be used to do the following:

- Access and restore individual files and folders from a virtual machine.
- Test a patch on an instantly recovered virtual machine before you apply the patch to production systems.
- Verify the backup image of the virtual machine and the applications.
- Verify an application within the instantly recovered virtual machine.
- Recover the instantly recovered virtual machine permanently by using Hyper-V Live migration or Storage migration. In a disaster recovery scenario, you can instantly recover a virtual machine in minutes and then schedule a migration to move it to a permanent storage on a Hyper-V host. The instantly recovered virtual machine remains available even during the migration process, which decreases the amount of downtime.

Note: You cannot back up instantly recovered virtual machines with the Agent for Hyper-V until you migrate the virtual machine from Backup Exec server storage and also remove from the virtual machine from Backup Exec server storage.

If you remove an instantly recovered virtual machine, any changes that you made are lost. Migrate the virtual machine from the Backup Exec server storage and remove it from Backup Exec to retain the changes or back up the instantly recovered virtual machine with the Agent for Hyper-V.

When you run an instant recovery job, the selected backup set is exposed to the Hyper-V host through an SMB share that is created on the Backup Exec server. The instantly recovered virtual machine disks are on the Backup Exec storage but they use the CPU of the Hyper-V host for their functions. All read operations are redirected to the Backup Exec server and the write operations are saved in a differencing disk at the location that is mentioned in the **Destination for VM registration and checkpoint** field when you create an instant recovery job. This path is on the Hyper-V host on which you want to recover the virtual machine.

Note: Backup Exec alerts you every week about the number of instantly recovered virtual machines running on the server. By default, the alert is triggered every Friday at 2.00 PM.

The following table describes the instant recovery process for a virtual machine.

Table D-5 Instant recovery process for a Hyper-V virtual machine

Step	Description
Step 1	You run an instant recovery job from a backup of a Hyper-V virtual machine.
Step 2	The Backup Exec server virtualizes the backup set.
Step 3	Backup Exec creates an SMB share.
Step 4	Backup Exec creates a virtual machine on the Hyper-V host.
Step 5	Backup Exec creates a snapshot of the virtual machine so that the writes can be made to the local disk.

Table D-5 Instant recovery process for a Hyper-V virtual machine *(continued)*

Step	Description
Step 6	Backup Exec starts up the virtual machine automatically if you select the option to power on the virtual machine after it is recovered.
Step 7 (optional)	You use Live Migration or Storage Migration to migrate the virtual machine from Backup Exec server storage if you want to save any changes that were made to the virtual machine.
Step 8	You run a job to do one of the following: ■ Remove an instantly recovered virtual machine that you no longer need to use. ■ Remove the instantly recovered virtual machine from Backup Exec server storage after you migrate the virtual machine. Warning: You cannot upgrade Backup Exec until you remove all of the instantly recovered virtual machines.

Instant Recovery Resiliency

As Backup Exec has enhanced the resiliency for instantly recovered virtual machines, if you restart the Backup Exec server, restart the Hyper-V server, or if there is a network connectivity issue, any changes that you made to the virtual machine are no longer lost. After either of these servers restart, the Backup Exec services start up and the virtualization process continues.

There are four scenarios for resiliency when the instantly recovered virtual machine is running on the Hyper-V host:

- Backup Exec server restarts and the Hyper-V server is running.
- Hyper-V server restarts and the Backup Exec server is running.
- Backup Exec and the Hyper-V servers restart.
- Network connectivity issues result in connection loss between the Backup Exec server and Hyper-V host.

In all these scenarios, the virtual machine starts automatically when the server restart is complete or network connectivity is restored. If the virtual machine does not start, you may require to restart it on the Hyper-V host.

Note: The virtual machine cannot be used until the server restart is complete or network connectivity is restored.

Instant Recovery Resiliency for Hyper-V uses the CORBA communication method. A configuration change may be required if you require to change the CORBA communication port on the Backup Exec server. The Backup Exec server and the Hyper-V server hosting the instantly recovered virtual machine must have the same CORBA port setting.

To change the CORBA port setting on the Backup Exec server

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Under **Custom port Number (Oracle only)**, select the **Use a custom port to receive operation requests from the Oracle server** check box and enter the same port number as entered on the Hyper-V server.
- 4 Stop and restart all Backup Exec services and rerun the backup.

To change the CORBA port setting on the Hyper-V host

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.
- 2 Click the **Database Access** tab.
- 3 Select the **Use a custom port to connect to the Backup Exec server during Oracle operations** check box.
- 4 Enter a port number that is not in use and can be used by Backup Exec and then click **OK**.

Note: This port number must match with what is configured on the Backup Exec server.

- 5 Restart the Backup Exec Remote Agent Service on the Hyper-V server.

Difference between an instantly recovered virtual machine and a restored virtual machine

Instant recovery of a virtual machine is different from a virtual machine restore in some aspects.

Table D-6 Differences between an instantly recovered virtual machine and a restored virtual machine

Item	Instant recovery of a virtual machine	Restore of a virtual machine
Data transfer	Does not transfer the virtual machine data to the instantly recovered virtual machine.	Transfers all data from the backup set to the restored virtual machine.
Job time	Instant recovery job runs instantly and no backup data is transferred. Therefore, the job time depends on the time taken to share the backup set and register the virtual machine.	Restore time depends on the size of the virtual machine and the network speed and storage speed.
Read/write operations	Uses the backup set image for all read operations. It uses a snapshot on the Hyper-V server for all write operations.	All data is already moved to the Hyper-V server. Therefore, there is no dependency on the Backup Exec server.
Data storage	Uses the Backup Exec server storage until you migrate the instantly recovered virtual machine.	Already uses the Hyper-V server storage.
Server restart	As the Instant Recovery resiliency is enhanced, if the Backup Exec server or the Hyper-V server restarts, the instantly recovered virtual machines remains accessible. In case of network connectivity issues the instantly recovered virtual machines are accessible after the connectivity is restored. If the virtual machine does not start, you may require to restart it on the Hyper-V host.	Restarting the Backup Exec server or the Hyper-V server has no effect on the restored virtual machine.

See [“Requirements for instant recovery of a Hyper-V virtual machine”](#) on page 1103.

See [“Creating an instant recovery job for a Hyper-V virtual machine”](#) on page 1104.

See [“About removing an instantly recovered Hyper-V virtual machine”](#) on page 1107.

See [“Notes about instant recovery of a Hyper-V virtual machine ”](#) on page 1103.

Requirements for instant recovery of a Hyper-V virtual machine

Before configuring an instant recovery job, review the following requirements:

- Ensure that sufficient disk space is available on the Hyper-V host to store all changes, such as virtual disk writes that happen on the instantly recovered virtual machine.
- Enable Backup Exec's Granular Recovery Technology to recover individual items from Microsoft applications in the Hyper-V backup job. If your virtual machine is not eligible for GRT, you can enable the **Enable Instant Recovery of all virtual operating system types, even those not eligible for GRT** option. Note that enabling this option will change the storage format of your next full backup to a format that is compatible with instant recovery.
See [“Using Granular Recovery Technology \(GRT\) with the Agent for Hyper-V”](#) on page 1087.
- While creating an instant recovery job, Backup Exec displays any backup sets that are eligible for Instant Recovery.
See [“Using Granular Recovery Technology \(GRT\) with the Agent for Hyper-V”](#) on page 1087.
- Run the instant recovery job for a virtual machine only if the virtual agent-based, GRT-enabled backup sets are stored on a disk storage device. Tape storage, cloud storage, and disk cartridge devices such as RDX are not supported.
- In an environment that has only IPv6 configured network cards, the Instant Recovery feature is supported if the Hyper-V server accesses the Backup Exec server using either the FQDN or the NetBIOS name.
- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List. You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility List.

Notes about instant recovery of a Hyper-V virtual machine

Review the following notes before configuring an instant recovery job.

- Backup Exec Management Command Line Interface (BEMCLI) is not supported for instant recovery jobs.

- Backup Exec does not support instant recovery of a virtual machine to a Hyper-V server that runs an older version than the Hyper-V server from which you backed up the virtual machine.
- The following apply in a CAS environment:
 - When backup sets are targeted to deduplication storage devices, the **Instant Recovery** option may be grayed out on Managed Backup Exec Server. It is recommended to run the instant recovery job on the server that has run the backup and duplicate job.
 - The central administration server or the managed Backup Exec server can instantly recover a virtual machine that is running on the server itself. However, only the managed Backup Exec server can remove the instantly recovered virtual machine that was recovered on the server.
 - You cannot administer an instant recovery job from a central administration server after you delegate the job to a managed Backup Exec server.
- In the following scenarios, Backup Exec stores disks in a different format than the disk format of the original virtual machine:
 - If the storage property of the disks that are attached to the original virtual machine was set to Fixed Size, Backup Exec converts the disk storage property to Dynamically Expanding during the backup process. When you create an instantly recovered virtual machine, the disk storage property is Dynamically Expanding.
- An instant recovery job does not preserve the saved state of the backup set that was used to create an instantly recovered virtual machine.
- An instant recovery job does not preserve the user-created snapshots that are present in the backup set used to create an instantly recovered virtual machine.
- Backup Exec supports instant recovery of a generation 2 virtual machine only if it was backed up on a Backup Exec server that was installed on Windows Server 2012 or later.

Creating an instant recovery job for a Hyper-V virtual machine

You can create an instant recovery job for a Hyper-V virtual machine, and then recover the virtual machine to the original location or to an alternate location.

Note: For applications such as SharePoint and Exchange, recover all of the virtual machines that the applications need to use to function properly. For example, to create a Microsoft Exchange environment, recover the virtual machines that run the Exchange client and the Active Directory, and then establish a connection between these two virtual machines.

See [“Requirements for instant recovery of a Hyper-V virtual machine”](#) on page 1103.

Complete the following steps to create an instant recovery job for a Hyper-V virtual machine:

To create an instant recovery job for a Hyper-V virtual machine

- 1 On the **Backup and Restore** tab, select the virtual machine that you want to instantly recover.
- 2 In the **Instant Recovery** group, click **Instantly Recover a VM**.
- 3 On the **Instantly Recover a VM** dialog box, in the **Job name** field, type a name for the instant recovery of a virtual machine job.
- 4 In the **Backup set selection** group box, in the **Show backup sets from** field, select the beginning and ending dates for the backup sets that you want to include in the backup set selection.

By default, only the backup sets for the jobs that ran in the past 30 days are displayed.

- 5 In the **Disk-based backup set** field, select the backup set that you want to use to create the instantly recovered virtual machine.

Only backup sets that meet the requirements for instant recovery are included in the list.

- 6 In the left pane, select **Destination**, and then set the following options for the job:

Item	Description
Virtual machine name	Enter the name that you want to appear in the Hyper-V manager console for this instantly recovered virtual machine.
Hyper-V server name	Select the name of the Hyper-V server where you want to instantly recover the virtual machine.
Server logon account	Select the name of the logon account that is required to access the server.
Destination for VM registration and checkpoint	Select the path on the Hyper-V host where you want all configuration details and the checkpoint file (differencing disk) of the recovered virtual machine to be stored.
Power on the virtual machine after it is recovered	Select this option to automatically start the virtual machine after it is instantly recovered.

- 7 On the **Instantly Recover a VM** dialog box, in the left pane, select **Schedule**, and then select the schedule of this job.
- 8 (Optional) On the **Instantly Recover a VM** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the job completes.
- 9 Click **OK**.

See [“Post-instant recovery tasks ”](#) on page 1106.

Post-instant recovery tasks

Perform these pos-instant recovery tasks before using the instantly recovered virtual machine.

- Backup Exec disables the network card on the instantly recovered virtual machine during the instant recovery process. To connect the computer to the network, you must configure the network settings on the computer.
- If you want to move the instantly recovered virtual machine to a Hyper-V host, you can migrate the virtual machine data files or the required disks from the instantly recovered virtual machine to the Hyper-V host using Hyper-V

Live Migration or Storage Migration. Note that during migration, the virtual machine data files are transferred to the host machine while the virtual machine is still running.

See the Microsoft website for the migration steps.

About removing an instantly recovered Hyper-V virtual machine

You should remove an instantly recovered virtual machine from the Backup Exec server storage after it has been migrated or when you no longer need to use it. Removing an instantly recovered virtual machine cleans all database entries and all folders that are created at the time of the instant recovery job.

If you remove the virtual machine from the Hyper-V host using the Hyper-V manager, you should still remove the instantly recovered virtual machine from the Backup Exec. When you remove an instantly recovered virtual machine from the Hyper-V host, the configuration and the checkpoint folders are deleted from the Hyper-V storage and the SMB share that was created for the instantly recovered virtual machine on the Backup Exec server is also removed.

When you run the **Remove a Recovered VM** job, Backup Exec checks the storage status of the disks, and then unregisters the virtual machine from the Hyper-V host.

The following disk status scenarios are possible:

- If all disks of the instantly recovered virtual machine are still running on the Backup Exec storage, Backup Exec first unregisters that virtual machine from the Hyper-V host and then proceeds with the removal process.
- If none of the virtual machines disks are running on the Backup Exec storage, Backup Exec proceeds with the removal process. The virtual machine is still available even after the remove operation and you can delete the virtual machine from the Hyper-V manager.
- If some of the disks are still running on the Backup Exec storage and some disks have been moved to a permanent storage, the **Remove a Recovered VM** job fails. You can either move all disks from the Backup Exec storage and then re-run the job or you can check the **Remove the virtual machine even though disks are on both Backup Exec server storage and the virtual machine host** check box if you do not want to save the changes.
- If the Hyper-V Replica feature is enabled on the instantly recovered virtual machine, the remove job fails. You can remove replication for this virtual machine and then run the job again.
- If the instantly recovered virtual machine migration is in progress, you can run the job again when the migration is over.

Note: Data lifecycle management (DLM) is postponed for the backup set that was used to create the instantly recovered virtual machine until you remove the virtual machine. The next cycle of DLM expires the backup set.

Removing an instantly recovered Hyper-V virtual machine

You should remove an instantly recovered virtual machine from the Backup Exec server storage after it has been migrated or when you no longer need to use it.

To remove an instantly recovered virtual machine

- 1 On the **Backup and Restore** tab, select the Hyper-V server that contains the instantly recovered virtual machine that you want to remove.

You can double-click a Hyper-V server and then in the left pane, click **Recovered VMs** to view the virtual machines that were instantly recovered to the server.

- 2 In the **Instant Recovery** group, click **Remove a Recovered VM**, and then do one of the following:

To remove a recovered virtual machine now

Do the following in the order listed:

- Select **Use Defaults and remove now**.
- On the **Select Recovered VMs to Remove** dialog box, select one or more recovered virtual machines to remove.
- Click **OK**.

To customize settings and then remove a recovered virtual machine Do the following in the order listed:

- Select **Customize settings and remove**.
- On the **Select Recovered VMs to Remove** dialog box, select one or more recovered virtual machines to schedule for removal.
- Click **OK**.
- On the **Remove a Recovered Virtual Machine** dialog box, in the **Job name** field, type a name for the job.
- In the **Server logon account** field, add or edit a logon account for the Hyper-V server.
- Check the **Remove the virtual machine even though disks are on both Backup Exec server storage and the virtual machine host** check box to remove the disks and then continue the job to remove the instantly recovered virtual machine. If you do not select this option, the disks cannot be removed and the job fails.
- On the **Remove a Recovered Virtual Machine** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.
- (Optional) On the **Remove a Recovered Virtual Machine** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the job completes.
- Click **OK**.

Best practices for instant recovery of a Hyper-V virtual machine

Best practices include tips and recommendations to help you use the instant recovery of a Hyper-V virtual machine feature effectively.

- Remove or migrate the instantly recovered virtual machines from Backup Exec before upgrading or uninstalling the Backup Exec server. Uninstalls are blocked on the Backup Exec servers and the server on which the Agent for

Hyper-V is installed if an instantly recovered virtual machine is running on these servers.

- If a large number of instantly recovered virtual machines are running at the same time, it may affect the performance of the Backup Exec server. You should periodically review the instantly recovered virtual machines that are running in your environment. Remove a virtual machine that you no longer need or migrate the virtual machine to the host and then remove the virtual machine from the Backup Exec server.

Migration uses network bandwidth. You should perform the migration during a time when the bandwidth requirement of other processes is low.

- When you migrate an instantly recovered virtual machine, you should use a path on the Hyper-V host that is different from the original location that the instantly recovered virtual machine was recovered.
- An instantly recovered virtual machine created on a Hyper-V host that has Backup Exec running in a virtual machine on the same host, must not have its destination set to a volume that also hosts the disks for the Backup Exec virtual machine.

About Recovery Ready for Hyper-V virtual machines

With the **Recovery Ready** feature, you can use the **Validate VM for Recovery** operation to validate the recoverability of virtual machines. When you create and run a Validate Virtual Machine for Recovery job, it runs tests on the virtual machine after which it is marked as recovery ready.

Recovery ready virtual machines can be used to do the following:

- Disaster recovery readiness: In a disaster recovery scenario, administrators are sure that the validated virtual machines are recoverable.
- Vault the backups to cloud or tape: Validates the backup sets before you vault them to devices such as cloud and tape.
- Audit and compliance of backups: For meeting company audit and regulatory compliance requirements, you can provide validation information of your backups for virtual machines.

When the Validate Virtual Machine for Recovery job runs, the virtual machine is registered on the Hyper-V server with the host name in the `Validate_VM name_GUID` format and then powered on. After the power-on is complete, a heartbeat check is run to check if the Hyper-V Heartbeat Service is running.

During virtual machine validation there is no data transfer.

When you create the Validate Virtual Machine for Recovery job, before the validation job runs, you can select the maximum time that is allowed for a virtual machine to boot. The default value is 10 minutes. You can select a value between 1 to 60 minutes.

See [“Configuring default backup settings”](#) on page 778.

After these checks are successfully completed, the virtual machine is marked as ready for recovery.

All information about the validation is part of the job log. After validation you can generate the **Recovery Ready Validation Summary** report to view the summary of the virtual machines that you validated.

See [“Recovery Ready Validation Summary”](#) on page 832.

To see the validation status, on the **Backup and Restore** tab, double-click the host name or on the **Storage** tab, double-click the disk storage name. In the left pane, click **Backup Sets**. For the selected server or storage view, the **Validation Status** is displayed.

The following table describes the process of validating Hyper-V virtual machines for recovery.

Table D-7 Virtual machine validation process

Step	Description
Step 1	Add the Hyper-V server.
Step 2	Take a disk-based GRT backup of the virtual machine that is hosted on the Hyper-V server. The Recovery Ready feature supports full, incremental, and deferential backup sets.
Step 3	Run a Validate Virtual Machine for Recovery job for a Hyper-V virtual machine.
Step 4	Backup Exec virtualizes the backup set.
Step 5	Backup Exec creates an SMB share.
Step 6	Backup Exec runs the following tests in the given order on the Hyper-V virtual machine being validated. ■ Register the virtual machine ■ Power On the virtual machine ■ Heartbeat check of the virtual machine

Table D-7 Virtual machine validation process (*continued*)

Step	Description
Step 7	If all the tests are successful, Backup Exec then marks the Hyper-V virtual machine as validated.

Tests run on a Hyper-V virtual machine being validated

The following tests are run on a Hyper-V virtual machine that is being validated.

Table D-8 Tests run on a Hyper-V virtual machine

Test	Description
Register the virtual machine	The first test is to register the virtual machine on the Hyper-V server.
Power ON the virtual machine	The second test is to power ON the registered virtual machine.
Heartbeat test	The third test is to check the Integration Services. If the services run successfully on the powered on virtual machine, the heartbeat test passes. Integration Services must be installed and the Hyper-V Heartbeat Service must be running for Hyper-V. These services must be running so that the heartbeat check is successful on the backed up virtual machine. If Integration Services are not installed, the validation job fails and for the corresponding backup set, Validation failed status is displayed.

All information that is related to the test results is also available in the job log.

Validation status for a Hyper-V virtual machine

During validation, the following statuses are displayed for the backup set of a Hyper-V virtual machine:

- **Validation successful:** The backup set has passed all tests and is recovery ready.
- **Validation failed:** The backup set has failed the tests and is not recovery ready.
- **Unable to validate:** The virtual machine is not validated because of environmental issues. Hence, the backup sets cannot be validated. If the Hyper-V Heartbeat Service is disabled at the time of backup, validation fails but the backup set is marked as **Unable to validate**.

During validation the following statuses are displayed for the validation job of a Hyper-V virtual machine:

- **Successful:** The validation job has passed all tests and the virtual machine is recovery ready.
- **Success with Exceptions:** The validation job has passed all tests but when there is a clean-up of resources that are used for validation, the power off fails.
- **Failed:** The validation job has failed the tests or validation is not performed.

If validation fails or Backup Exec is unable to validate a virtual machine for recovery, you can check the validation job log for more details.

The following matrix gives details about the tests and validation status.

Table D-9 Validation status for a Hyper-V virtual machine

Validation test	Backup set status	Validation job status	Further steps/Reason for failure
Register virtual machine	If registration is successful, go to the next check.		
	Validation failed	Failed	Backup set issue
	Unable to validate	Failed	Environment issue
Power ON virtual machine	If power ON is successful, go to the next check.		
	Unable to validate	Failed	Connectivity issue or check not attempted as registration failed
Heartbeat check	Unable to validate	Failed	Test not attempted
	Validation failed	Failed	Boot issue or Integration Services issue
	Validation successful	Successful	Virtual machine is validated and recovery ready

See [“Requirements for validating a Hyper-V virtual machine for recovery ”](#) on page 1114.

See [“Notes about validating a Hyper-V virtual machine for recovery ”](#) on page 1114.

See [“Creating a validate virtual machine for recovery ”](#) on page 1115.

Requirements for validating a Hyper-V virtual machine for recovery

Review the following requirements before you configure a Validate virtual machine for Recovery job for a Hyper-V virtual machine:

- Hyper-V Integration Services must be installed on the Hyper-V server.
- Ensure that sufficient disk space is available on the Hyper-V host to store all changes, such as virtual disk writes on the validated virtual machine.
- While creating a Validate Virtual Machine for Recovery job, Backup Exec displays any backup sets that are eligible for Instant Recovery..
- Run the Validate Virtual Machine for Recovery job for a virtual machine only if the virtual agent-based backup sets are stored on a disk storage device. Tape storage, cloud storage, and disk cartridge devices such as RDX are not supported.
See [“Using Granular Recovery Technology \(GRT\) with the Agent for Hyper-V”](#) on page 1087.
- In an environment that has only IPv6 configured network cards, the Validate VM for Recovery operation is supported if the Hyper-V server accesses the Media server using either the FQDN or the NetBIOS name.
- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List. You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

Notes about validating a Hyper-V virtual machine for recovery

Review the following notes before configuring a Validate Virtual Machine for Recovery job for a Hyper-V virtual machine:

- Backup Exec Management Command Line Interface (BEMCLI) is not supported for validate virtual machine for recovery jobs.
- Backup Exec does not support validate virtual machine for recovery to a Hyper-V server that runs an older version than the Hyper-V server from which you backed up the virtual machine.
- Backup Exec supports validate virtual machine for recovery of a generation 2 virtual machine only if it was backed up with Backup Exec installed on Windows 2012 or later.
- The following applies in a CAS environment:
 - When backup sets are targeted to deduplication storage devices, the **Recovery Ready** option may be grayed out on Managed Backup Exec

Server. It is recommended to run the recovery ready job on the server that has run the backup and duplicate job.

- If the Central Administration Server (CAS) has a backup set, create the Validate Virtual Machine for recovery job only on the CAS server. If the Managed Backup Exec Server (MBES) has a backup set, create the validate virtual machine for recovery job only on the MBES server.
- You cannot administer a validate virtual machine for recovery job from a central administration server after you delegate the job to a managed Backup Exec server.

Best practices about validating a Hyper-V virtual machine for recovery

Review the following best practices before configuring a Validate Virtual Machine for Recovery job for a Hyper-V virtual machine:

- While validating a virtual machine for recovery, select a destination folder on a volume that free space. The temporary virtual machine requires more disk space than the size of RAM configured for the virtual machine being validated for recovery.
- It is recommended that you have the latest Hyper-V integration services on the backed-up Hyper-V virtual machine.
- While validating a virtual machine for recovery, it is recommended that the destination Hyper-V server version be same as the Hyper-V server on which the virtual machine is hosted (when the backup was taken).

Creating a validate virtual machine for recovery

You can create a validate virtual machine recovery job for a Hyper-V virtual machine.

See [“Requirements for validating a Hyper-V virtual machine for recovery ”](#) on page 1114.

Complete the following steps to create validate virtual machine recovery job for a Hyper-V virtual machine.

To create a Validate Virtual Machine Recovery job

- 1 On the **Backup and Restore** tab, select the virtual machine that you want to validate.
- 2 In the **Recovery Ready** group, click **Validate VM for Recovery**.

- 3
- On the **Validate Virtual Machine for Recovery** dialog box, in the **Job name** field, type a name for the validation job or use the default name.
- 4
- In the **Backup set selection** group box, select the backup sets that you want to include in the backup set selection.

Item	Description
Use the latest available disk-based backup set when the job is run	Indicates that only the latest available disk-based backup set is used when the VMware virtual machine validation recovery job is run.
Select a disk-based backup set	Indicates that an available disk-based backup set is used for the validate Hyper-V virtual machine for recovery job.
	Show backup sets from
	Indicates the beginning and ending dates for the backup sets that you want to include in the backup set selection.
	By default, only the backup sets for jobs that ran in the past 30 days are displayed.
Disk-based backup set	Indicates the backup set that you want to use for the validation job.
	Only backup sets that meet the requirements for validation are included in the list.

- 5
- On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Destination**, and then set the following options for the job.

Item	Description
Hyper-V server name	Indicates the name of the Hyper-V server where you want to validate the virtual machine. Note: You can also select a Hyper-V server that is different from the source Hyper-V server.
Server logon account	Indicates the name of the logon account that is required to access the Hyper-V server.
Destination for VM registration and checkpoint	Browse a path on the Hyper-V host where all configuration details and the checkpoint file (differencing disk) of the validated virtual machine are stored.

- 6 On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.

Item	Description
Recurrence	Select this option to create a recurring schedule for the job.
Run now with no recurring schedule	Runs the job immediately without a recurring schedule.
Run on	Schedules the job to run on a specific date and time.
Create without a schedule	Creates the job without scheduling it. When you use this option, the job does not run at the time of creation. The job remains unscheduled until you choose to run it.
Submit job on hold	Enables you to submit the job, but not run the job until you change the job's hold status.

- 7 (Optional) On the **Validate Virtual Machine for Recovery** dialog box, in the left pane, select **Notification**, and then select recipients to notify when the Validate a Virtual Machine for Recovery job completes.
- 8 Click **OK**.

The Validate a Virtual Machine for Recovery job runs and after it is completed successfully, the virtual machine is recovery ready.

To see the validation status, on the **Backup and Restore** tab, double-click the host name or on the **Storage** tab, double-click the disk storage name. In the left pane, click **Backup Sets**. For the selected server or storage view, the **Validation Status** is displayed.

Troubleshooting issues with Backup Exec Agent for Microsoft Hyper-V

This section contains troubleshooting strategies that can help resolve issues with Backup Exec Agent for Microsoft Hyper-V.

- Snapshot failures may occur during backup of a Linux virtual machine, if the integration services are not running inside the virtual machine. Refer to Microsoft documentation and ensure the latest integration services are deployed and running correctly inside the virtual machine.

- Access denied errors may occur when doing backup browse or backup of a virtual machine hosted on an SMB or Scale Out File Server share. This may be because the logon account provided in Backup Exec does not have access to the share.

Backup Exec Agent for Microsoft SQL Server

This appendix includes the following topics:

- [About the Agent for Microsoft SQL Server](#)
- [Requirements for using the SQL Agent](#)
- [About installing the SQL Agent](#)
- [Backup strategies for SQL](#)
- [Adding SQL Servers to the list of servers on the Backup and Restore tab](#)
- [Configuring Backup Exec to run a consistency check before every SQL backup](#)
- [Using snapshot technology with the SQL Agent](#)
- [Using database snapshots for SQL Server](#)
- [Backing up SQL databases and transaction logs](#)
- [Restoring SQL databases and transaction logs](#)
- [Disaster recovery of a SQL Server](#)
- [About SQL Server Always On availability groups](#)

About the Agent for Microsoft SQL Server

The Agent for Microsoft SQL Server (SQL Agent) enables network administrators to perform backup and restore operations on installations of SQL that are connected to a network. SQL database backups can be integrated with network backups without separate administration or dedicated hardware.

The SQL Agent provides support for the following:

- Database, transaction log, and differential backups, as well as database recovery and replacement. The support includes databases with In-Memory optimized tables.
- An automated restore of the system databases.
- Simplified Disaster Recovery, which automates the disaster recovery process of SQL Servers.
- Restores of SQL databases to alternate locations.
- Hot backup copies of SQL databases during backup operations. This feature enables you to direct a copy of the actual data streams being sent to media by a SQL database to a local directory for later use.
- Backups of multiple instances.
- Standby database. If the primary SQL Server fails, or is shut down for maintenance, another database called a standby database can be brought online.
- Database Consistency Checks (DBCC) for each backup and restore job, including a fast database consistency check of only the physical consistency of the database.
- Full, bulk-logged, and simple recovery models. With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run. Therefore, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.
- Restores of transaction logs to a specific point in time or to a named transaction when log marks are used.
- Database snapshots.
- Maintenance replication settings during redirected restores.
- A Verify Only option for a restore job that determines both the validity of the SQL data on the media and the ability of the destination SQL database to accept this data before the database is deleted or overwritten during a restore job.
- Back up with checksum generation. This option is used as a redundancy check and works with the Verify Only option on a restore job.
- Continuation of restore jobs when errors are detected. This feature enables you to restore as much data as possible from a corrupt database backup.

- Copy-only one-time backups, which enable you to copy a database without affecting the full-differential-log restore sequence.
- In SQL Server 2008 or later editions that support compression, you can use SQL software compression for backup jobs.

For information about the best practices to use Backup Exec Agent for Microsoft SQL Server (SQL Agent), refer to *Backup Exec Best Practices*.

See [“About installing the SQL Agent”](#) on page 1121.

Requirements for using the SQL Agent

The following are required for the SQL Agent:

- The Agent for Windows must be installed on any remote SQL Server that you want to back up.
- Backup Exec Logon account associated with the SQL server for backup and restore jobs must be:
 - A member of the local user group named ‘Administrators’
 - A member of SQL Server role named ‘sysadmin’

Note: SQL Server credentials are not supported.

For specific operating system requirements and supported SQL Server service packs for the SQL Agent, refer to the Backup Exec Software Compatibility List.

See [“Testing logon accounts”](#) on page 770.

See [“Backup Exec logon accounts”](#) on page 759.

About installing the SQL Agent

The SQL agent is installed as part of the Agent for Applications and Databases feature and can protect local or remote SQL server databases.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

Backup strategies for SQL

Backup Exec incorporates online, nondisruptive SQL database protection as part of everyday backup routines, which increases the chance of data recovery and

minimizes data loss without inhibiting daily database activity. Using database, differential, and log backups provides a good balance between backup windows and minimizes the amount of time that will be spent recovering a database if the need arises.

To decide which backup methods to use for the best data protection, consider the following for typical environments:

- In small environments, consider running a daily full database backup every evening and daily transaction log backups.
- In mid-sized environments, consider running a weekly full database backup and daily transaction log backups along with daily differential backups except on the day when the full backup is run.
- In large environments, consider running daily differential database backups, weekly full database backups, and transaction log backups as necessary. Many shops run full backups on a weekly basis, preferring to run differential backups throughout the week to keep backup run time to a minimum.

The trade-off with running fewer full backups and running more differential backups occurs at recovery time when you must recover using the full database backup as well as the last differential database backup, and all log backups made after the last differential database backup.

What will work best for you will be based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

When you develop a SQL backup strategy, consider the following:

Table E-1 Recommendations for backing up SQL

SQL Server backup strategies	Description
Protect the entire SQL Server.	To ensure SQL is completely protected, back up the following on a regular basis: ■ The system drive that SQL is on. ■ The Windows registry and System State. ■ Transaction logs.
When you upgrade, run new full database backups.	If you upgrade SQL, run new full database backups. You may not be able to restore backups from one version or service pack level of SQL to other versions.

Table E-1 Recommendations for backing up SQL *(continued)*

SQL Server backup strategies	Description
Run consistency checks before backups.	It is recommended that you run a consistency check before a backup. If a database or transaction log contains errors when it is backed up, the backup will still contain the errors when it is restored, if it is restorable at all. See “Configuring Backup Exec to run a consistency check before every SQL backup” on page 1124.
Back up your system databases regularly.	Back up the master database and service packs that are installed whenever procedures are run that change information in the database, especially after the following: ■ New databases are created. ■ Files are added to an existing database. ■ Usernames or passwords are added or changed. If changes are not backed up before the master database must be restored, the changes are lost.
Run one backup at a time.	Do not schedule more than one backup to occur simultaneously against a database or its transaction log.
Back up transaction logs on databases that are configured for the full recovery model.	Back up the transaction logs on databases because the transaction logs continue to grow if you do not back them up.

Adding SQL Servers to the list of servers on the Backup and Restore tab

You can add SQL Servers to the list of servers on the **Backup and Restore** tab so that you can select the SQL databases for backup.

To add a SQL Server to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Windows computers and servers**, and then click **Next**.
- 3 Follow the **Add a Server** wizard prompts to add the SQL Server to the list of servers in the **Backup and Restore** tab.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

Configuring Backup Exec to run a consistency check before every SQL backup

If you back up a database or transaction log that contains errors, these errors will still exist when the backup is restored. In some cases, this can prevent a successful restore. Backup Exec enables you to check the logical and physical consistency of the data before and after a backup. SQL reports any consistency check failures in the Backup Exec job log. It is strongly recommended that you always run a consistency check before the backup.

By default, the backup job default option **Consistency check before backup** is enabled for **Physical check only**.

Backup Exec’s consistency check uses the following SQL consistency check utilities:

- CHECKDB
- CHECKCATALOG
- PHYSICAL_ONLY

CHECKDB, CHECKCATALOG, and PHYSICAL_ONLY are performed for database-related operations.

For more information concerning these utilities, see your Microsoft SQL documentation.

To run a consistency check before a SQL backup

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select a backup job type.

For example, if you want to set up the default options for SQL Server backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices that you configure. You can configure different default options for backup jobs that you send to different types of storage.
- 3 On the **Backup Job Defaults** dialog box, on the left pane, click **Microsoft SQL**.
- 4 Ensure that a consistency check is enabled in the field **Consistency check before backup**.
- 5 Click **OK**.

See [“Backing up SQL databases and transaction logs”](#) on page 1127.

Using snapshot technology with the SQL Agent

Backup Exec uses snapshot technology by default for SQL Server backups. The SQL Agent supports full snapshot backups using Microsoft’s Volume Shadow Copy Service (VSS). The use of snapshot technology can reduce both restore time and backup performance on the server.

When you submit a backup job that uses snapshot technology, a snapshot of each volume is created, providing a point-in-time record of the data. Backup Exec uses snapshot technology to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. The data is then backed up from the snapshots, and then the snapshots are deleted.

Note: Use snapshot technology with jobs that use deduplication devices.

Before you use snapshot technology with the SQL Agent, review the following information:

- With snapshot technology, a point-in-time view of the SQL database is "snapped" and then backed up, leaving the actual SQL database open and available for users.
- SQL backups that use snapshot technology are considerably larger than regular SQL backups (also known as streaming backups).
- Performing consistency checks before backup is strongly recommended.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 1124.

- The SQL Agent supports only full snapshot backups; transaction log snapshots and differential snapshots are not supported.
- With the SQL Agent, snapshot and streaming backups are interoperable when you restore SQL data.
- Performing database consistency checks both before and after backups affects the time required for the backup jobs.

The following SQL backup options are not supported with snapshot backups:

- **Use checksums on backup (SQL 2005 or later)**
This option is used as a redundancy check, and works with the **Run verify only and do not restore data** restore option.
- **SQL Server 2008 Enterprise Edition software compression**
- **Create on-disk copies of SQL backups to be placed on the SQL Server where the database is located**

Note: Microsoft Virtual Shadow Copy Service (VSS) snapshots are not the same as SQL database snapshots. VSS snapshots enable you to create point-in-time snapshots of disk volumes and shares; database snapshots enable you to create point-in-time copies of SQL databases.

See [“Using database snapshots for SQL Server”](#) on page 1126.

Using database snapshots for SQL Server

SQL database snapshots enable you to quickly revert a database back to the state it was in when the database snapshot was created. When you use a database snapshot, a full restore of the host database is not required to revert the database. However, the changes that are made to the host between the time a database snapshot is created and the point at which it is reverted are lost.

The Backup Exec SQL Agent works with the SQL database to create database snapshots, which are read-only, point-in-time copies of an existing host database. When Backup Exec runs a SQL backup job using the Database Snapshot backup method, a request is sent to the host database instructing it to create a database snapshot.

Note: The snapshot backup method for SQL databases is only supported by SQL Server Enterprise Edition.

Database snapshots cannot be backed up to storage media. Rather, they are written to a SQL snapshot file on disk. After running the database snapshot job, Backup Exec creates history and job log information to indicate the job's status.

Because database snapshots cannot be backed up, all database snapshots will be lost if the disk where the host database is installed fails. Therefore, database snapshots should not be used as your sole database protection strategy. They should be used in conjunction with an overall Backup Exec database protection strategy that includes full, differential, and transaction log backups of the SQL database.

For more information, see your Microsoft SQL documentation.

Note: SQL database snapshots are not the same as Microsoft Virtual Shadow Copy Service (VSS) snapshots. Whereas VSS snapshots enable you to create point-in-time snapshots of disk volumes and shares, database snapshots enable you to create point-in-time copies of SQL databases.

Note: SQL database snapshot catalog information that refers to deleted database snapshots is periodically removed from the catalogs. If backup media is re-cataloged, the database snapshot catalog information will be periodically removed again.

SQL Server database snapshots created with Backup Exec can be used to revert a SQL database back to the state it was in at a previous point in time, without having to run a full database restore job.

When you view SQL database snapshots in the resource view in the Restore Wizard, they appear as backup sets, in chronological order with the most recent snapshot appearing first.

The following caveats apply when reverting a database:

- You cannot undo a SQL database that has been reverted.
- Before reverting a database, Backup Exec deletes all existing database snapshots with the exception of the snapshot used for the revert. After being deleted, the database snapshots cannot be recovered.
- You cannot redirect a database snapshot restore job.

Backing up SQL databases and transaction logs

Backup Exec includes three methods for backing up databases: Full, Differential, and Full Copy-only. The full method backs up the entire database including all

system tables. The differential method backs up only the changes made to the database since the last full backup. The copy method works in the same manner as the full method, except that it does not affect future differential or log backups.

A differential backup is smaller and faster than a full backup, so differential backups can be run more often than full backups. Because differential backups allow the restore of a system only to the point that the differential backup was created, you should also create multiple log backups between the differential backups. Using transaction log backups allows you to recover the database to the exact point of failure.

Consider using differential backups when only a relatively small amount of data changes between full backups, or if the same data changes often. Differential backups may also work well in your environment if you are using the simple recovery model and need backups more often, but cannot spare the time to do frequent full backups. If you are using the full or bulk-logged recovery models, you can use differential backups to decrease the time it takes to roll forward log backups when restoring a database.

If you want to run database backups only, instead of a mix of database and log backups, use the simple recovery model for the database so that the transaction log is automatically truncated when a checkpoint occurs in the database. This helps prevent transaction logs from becoming full since with other recovery models the logs are not cleared after a database backup.

With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run.

If you do not run transaction log backups, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.

System databases can only be backed up with the full method; you cannot use the log or differential methods to back up the master database.

Note: You cannot back up databases to storage that is attached to a computer on which the Remote Media Agent for Linux Servers is installed.

The SQL Agent supports a mirrored SQL database configuration, although Microsoft places the following limitations on the mirroring of SQL databases:

- You cannot back up or restore a mirrored SQL database. If you attempt to back up or restore a mirrored database, the backup job or restore job fails.
- You cannot restore the primary SQL database while it is configured in a mirrored configuration. To restore the primary SQL database, you must stop database mirroring of the primary database.

- You can back up a primary SQL database and its transaction logs only if the backup job does not leave the database in a non-recovered state.

You can set backup job default options for all SQL backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Automatic exclusion of SQL data during volume-level backup

If you select a volume that contains SQL data for backup, the SQL Agent determines which SQL data should not be included in a volume-level backup. For example, .MDF and .LDF files should not be part of the backup because they are opened for exclusive use by the SQL system. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

While it is not recommended, if you want to include SQL data in a volume-level backup, you must first dismount the database that you want to back up. Then, run the backup job.

How to back up SQL clusters

You may need to manually add resource containers for SQL clusters before you can back up the databases.

To add resource containers, install the Agent for Windows on the cluster physical nodes. If the resource container for the virtual SQL server is not automatically detected, use the **Add Server** wizard to add the virtual resource container for the SQL virtual cluster node. When you run the **Add Server** wizard, uncheck the **Upgrade the Agent for Windows** option because it is already installed on the physical node. Then, make your backup selections from the virtual resource container that you added.

How to back up SQL transaction logs

When you run log backups, you should use Backup Exec exclusively to perform log transaction backups.

Backup Exec provides the Log and Log No Truncate methods for backing up transaction logs.

Use the Log No Truncate method only when the database is corrupted or if database files are missing. This method backs up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the last database backup and any

previous transaction log backups to restore the database to the point at which it failed. However, any uncommitted transactions are rolled back. The Log No Truncate method does not remove committed transactions after the log is backed up.

To use the Log No Truncate backup to restore a database, you should also have a database backup that was created before the Log No Truncate backup. The transaction log contains only the log files that are used in the restore process, which alone are not sufficient to restore a complete database. You must have at least one database backup and a log backup of the database to restore the database.

Caution: Do not run a log backup using either method if the SQL database uses the simple recovery model. With the simple recovery model, you can recover data only up to the most recent full or differential backup. If you run a log backup on a database using the simple recovery completion state, the backup completes with exceptions.

To check the database properties, from the Database management tools on the SQL Server, right-click the database, click **Properties**, click the **Options** tab, and then view the configuration settings.

To back up SQL databases and transaction logs

- 1 On the **Backup and Restore** tab, right-click a SQL Server, and then right-click the selection.

To select multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.

Note: You can select the SQL databases to back up on the **Browse** tab. In the right pane of the **Backup Selections** dialog box, you can view the Name, Size, Type, Modified Time, and Attributes for the selection. The Attributes provide the status of the database, so if there are any problems you can resolve them before you run the backup job. You can also include or exclude specific files or specific types of files using the **Selection Details** tab.

- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, select the schedule for this job.
- 8 On the **Backup Options** dialog box, select the storage device for this job.
- 9 On the **Backup Options** dialog box, in the left pane, select **Microsoft SQL**.
- 10 Set any of the following options for this job:

Backup method

Select one of the following backup methods that you want to use for this job:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

Backup method

Select one of the following SQL-specific backup methods that you want to use for this job:

■ **Full - Back up databases**

This option backs up the entire database. This option is selected by default.

■ **Full Copy - Back up databases (copy)**

This option backs up the entire database without affecting future differential or log backups.

Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.

After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set that is required to run future differential backups.

■ **Automatic - Backup up transaction log if enabled and then back up database changes since the last full or incremental**

This option lets you back up the entire SQL instance even though some databases may not support log backups. All of the databases are backed up using the Incremental (block level) backup method. In addition, the databases that support log backups are backed up using the Log backup method.

Note: If snapshot is not enabled, an Incremental (block level) backup method cannot be performed and the Differential backup method is used.

■ **Log - Back up and truncate transaction log**

This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).

If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.

Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 1124.

- **Log No Truncate - Back up without truncating transaction log**

This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.

- **Differential - Backup up database changes since the last full**

This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point in time that the differential backup was created, you should also create multiple log backups between the differential backups.

- **Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job**

This option backs up all of the blocks of data and logs that have been created or modified since the last full backup.

- **Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job**

This option backs up all of the blocks of data and logs that have been created or modified since the last full or incremental backup.

- **Database Snapshot - Read-only point-in-time copy of databases**

This option creates a read only, point-in-time copy of another database.

See ["Using database snapshots for SQL Server"](#) on page 1126.

Note: SQL differential or incremental backups are supported for conversion to virtual when you use the Automatic, Differential (block-level), or Incremental (block-level) backup methods.

Note: SQL Full and differential backups always run on the primary node of the availability group. Other backup methods use the selection made in the Backup Preferences setting of the availability group.

Additionally, Backup Exec runs a full backup when you select the Automatic or Log backup methods if a full backup was not previously run on the database. A full backup also runs for one of the following conditions:

- A new database is added or restored.
- Backup Exec did not run the last full backup.
- Only Full Copy and Incremental backups were run on the database instead of Full backups.

See [“Configuring backup methods for backup jobs”](#) on page 197.

Consistency check before backup

Select one of the following consistency checks to run before a backup:

- None.
This option does not run a consistency check before a backup. It is recommended that you always run a consistency check either before or after the backup.
- Full check, excluding indexes.
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- Full check, including indexes.
This option includes indexes in the consistency check. Any errors are logged.
- Physical check only.
This option performs a low-overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 1124.

Continue with backup if consistency check fails

Choose if you want to continue with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.

Consistency check after backup

Select the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure that the data was consistent at the time of the backup

The following options are available:

- **None.**
This option does not run a consistency check after a backup. It is recommended that you always run a consistency check after the backup. This option is selected by default.
- **Full check, excluding indexes.**
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- **Full check, including indexes.**
This option includes indexes in the consistency check. Any errors are logged.
- **Physical check only.**
This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

Use checksums on backup (SQL 2005 or later)

Choose to add the checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the restore option **Run verify only and do not restore data**. Use this option and the **Run verify only and do not restore data** option to ensure that you restore from a verified SQL backup during a restore of the SQL database.

Create on-disk copies of SQL backups to be placed on the SQL server where the database is located

Choose to create an on-disk copy of the SQL database that you want to back up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the **Save to path** option.

This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores.

Note: This option does not support snapshot technology.

Save to path

Specify a path in which to save on-disk copies of SQL backups.

SQL Server 2008 Enterprise Edition software compression

Select a compression setting that you want to use for this backup job:

- **None.**
Do not use compression.
- **Compress.**
Use SQL Server 2008 or later compression if it is supported by the SQL Server instance that is installed.

SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.

You can find a list of compatible operating systems, platforms, applications, and supported service packs in the Backup Exec Software Compatibility List.

It is recommended that you do not use SQL 2008 or later software compression in a backup job that also uses Backup Exec-initiated software compression. Minimal benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.

SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses Advanced Open File options.

Note: You cannot use this option for backup jobs that deduplicate data.

One-time backup method Specifies one of the following methods for one-time backups:

■ **Full - Back up databases**

This option backs up the entire database. This option is selected by default.

See [“Backing up SQL databases and transaction logs”](#) on page 1127.

■ **Full Copy - Back up databases (copy)**

This option backs up the entire database without affecting future differential or log backups.

Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.

After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

11 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.

12 Click **OK**.

13 On the **Backup Definition Properties** dialog box, click **OK**.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Editing backup definitions”](#) on page 205.

See [“Changing default backup job settings”](#) on page 640.

Setting default backup options for SQL Server

You can use the defaults that Backup Exec sets during installation for all SQL Server jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for SQL Server

1 Click the Backup Exec button, and then select **Configuration and Settings**.

2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for SQL Server backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for backup jobs that you send to different types of storage.

- 3 In the left pane, select **Microsoft SQL**.
- 4 Select the appropriate options.

Backup method

Select one of the following backup methods that you want to use for this job:

■ Full - Back up databases

This option backs up the entire database. This option is selected by default.

■ Full Copy - Back up databases (copy)

This option backs up the entire database without affecting future differential or log backups.

Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.

After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

Backup method

Select one of the following SQL-specific backup methods that you want to use for this job:

■ **Full - Back up databases**

This option backs up the entire database. This option is selected by default.

■ **Full Copy - Back up databases (copy)**

This option backs up the entire database without affecting future differential or log backups.

Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.

After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set that is required to run future differential backups.

■ **Automatic - Backup up transaction log if enabled and then back up database changes since the last full or incremental**

This option lets you back up the entire SQL instance even though some databases may not support log backups. All of the databases are backed up using the Incremental (block level) backup method. In addition, the databases that support log backups are backed up using the Log backup method.

Note: If snapshot is not enabled, an Incremental (block level) backup method cannot be performed and the Differential backup method is used.

■ **Log - Back up and truncate transaction log**

This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).

If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.

Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 1124.

- **Log No Truncate - Back up without truncating transaction log**

This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.

- **Differential - Backup up database changes since the last full**

This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point in time that the differential backup was created, you should also create multiple log backups between the differential backups.

- **Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job**

This option backs up all of the blocks of data and logs that have been created or modified since the last full backup.

- **Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job**

This option backs up all of the blocks of data and logs that have been created or modified since the last full or incremental backup.

- **Database Snapshot - Read-only point-in-time copy of databases**

This option creates a read only, point-in-time copy of another database.

See ["Using database snapshots for SQL Server"](#) on page 1126.

Note: SQL differential or incremental backups are supported for conversion to virtual when you use the Automatic, Differential (block-level), or Incremental (block-level) backup methods.

Additionally, Backup Exec runs a full backup when you select the Automatic or Log backup methods if a full backup was not previously run on the database. A full backup also runs for one of the following conditions:

- A new database is added or restored.
- Backup Exec did not run the last full backup.
- Only Full Copy and Incremental backups were run on the database instead of Full backups.

See [“Configuring backup methods for backup jobs”](#) on page 197.

Consistency check before backup

Select one of the following consistency checks to run before a backup:

- None.
This option does not run a consistency check before a backup. It is recommended that you always run a consistency check either before or after the backup.
- Full check, excluding indexes.
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- Full check, including indexes.
This option includes indexes in the consistency check. Any errors are logged.
- Physical check only.
This option performs a low-overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 1124.

Continue with backup if consistency check fails

Choose if you want to continue with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.

Consistency check after backup

Select the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure that the data was consistent at the time of the backup.

The following options are available:

- **None.**
This option does not run a consistency check after a backup. It is recommended that you always run a consistency check after the backup. This option is selected by default.
- **Full check, excluding indexes.**
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- **Full check, including indexes.**
This option includes indexes in the consistency check. Any errors are logged.
- **Physical check only.**
This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

Use checksums on backup (SQL 2005 or later)

Choose to add the checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the restore option **Run verify only and do not restore data**. Use this option and the **Run verify only and do not restore data** option to ensure that you restore from a verified SQL backup during a restore of the SQL database.

Create on-disk copies of SQL backups to be placed on the SQL server where the database is located

Choose to create an on-disk copy of the SQL database that you want to back up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the **Save to path** option.

This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores.

Note: This option does not support snapshot technology.

Save to path

Specify a path in which to save on-disk copies of SQL backups.

**SQL Server 2008
Enterprise Edition
software compression**

Select a compression setting that you want to use for this backup job:

- None.
Do not use compression.
- Compress.
Use SQL Server 2008 or later compression if it is supported by the SQL Server instance that is installed.

SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.

You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

It is recommended that you do not use SQL 2008 or later software compression in a backup job that also uses Backup Exec-initiated software compression. Minimal benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.

SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses Advanced Open File options.

Note: You cannot use this option for backup jobs that deduplicate data.

One-time backup method Specifies one of the following methods for one-time backups:

- Full - Back up databases

This option backs up the entire database. This option is selected by default.

See [“Backing up SQL databases and transaction logs”](#) on page 1127.

- Full Copy - Back up databases (copy)

This option backs up the entire database without affecting future differential or log backups.

Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.

After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

5 Click **OK**.

See [“Backing up SQL databases and transaction logs”](#) on page 1127.

Restoring SQL databases and transaction logs

The SQL Agent lets you restore SQL Server databases. You can restore the databases to their original location or you can redirect the restore to a new location. The number of jobs you decide on depends on the types of backup jobs that protect the database. If you use one job to restore a database, select all the backup sets that you want to apply. Include the full backup, any differential backups, and any log backups.

With very large databases this process can take several hours to complete. During this time Backup Exec reports that no data is being transferred, and the **Byte count** field in the Job Monitor view is not updated. When SQL has completed filling the files with zeros, the restore job continues. This occurs for all database restores but is noticeable only on very large databases.

Restoring SQL databases that contain In-Memory optimized tables

Ensure the target SQL server has adequate available memory when restoring SQL databases that contain In-Memory optimized tables. Refer to the Microsoft SQL server documentation for more information.

Restoring encrypted SQL databases

SQL 2008 supports Transparent Database Encryption (TDE), which lets you encrypt SQL 2008 databases at the backup set level.

When you back up a database that uses TDE, Microsoft recommends that you back up the certificate keys and encryption keys with the database. If you do not include the certificate keys and encryption keys, you must perform all backup and restore operations within the selected SQL instance.

Note: Backup Exec can redirect the restore of the database data that used TDE only if the certificate keys and encryption keys are applied to the destination instance. If the certificate keys and encryption keys are not applied to the destination instance, an error appears stating that the certificate thumbprint cannot be found. See your Microsoft SQL 2008 documentation.

How to restore from SQL transaction logs up to a point in time

You can restore transactions from a transaction log up to and including a point in time in the transaction log. After the point in time is reached, recovery from the transaction log is stopped. To find dates and times of transactions, check your client application event log.

If the specified point in time is later than the time contained in the most recent transaction log being restored, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state. If the specified point in time is before the time contained in the transaction log or logs being restored, no transactions are restored.

How to restore from SQL transaction logs up to a named transaction

You can restore transactions from a transaction log up to and including a named transaction (or mark). After the named transaction is reached, recovery from the transaction log is stopped.

Since named transactions do not necessarily have unique names, you can also specify a date and time after which the restore operation is to search for the named transaction. For example, if you specify a restore from a log up to the named transaction AfternoonBreak, found after 6/02/2000, 12:01 p.m., then the restore operation will not search for AfternoonBreak until after that time. To find dates and times of named transactions, check your client application event log.

If the named transaction is not found, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state.

Note: The names of transactions are case-sensitive. Ensure you enter the correct upper- and lower-case characters when specifying a named transaction.

Redirecting restores of SQL

You can redirect the following:

- A database backup to a different server, database, or instance.
- Differential and log backups to wherever the associated database is restored.
- A database from a 32-bit or 64-bit platform to any other platform.

You can use both single-job restores and multiple-job restores in redirected restore operations.

To restore SQL databases and transaction logs

- 1 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 2 Select **Microsoft SQL Server databases**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Restoring the SQL master database”](#) on page 1148.

Restoring the SQL master database

If the master database is damaged, symptoms may include the following:

- An inability to start SQL.
- Segmentation faults or input/output errors.
- A report generated by SQL Database Consistency Checker utility (DBCC).

If the master database is critically damaged and SQL cannot be started, rather than running the Rebuild Master utility or reinstalling SQL to be able to restart SQL, you can replace the corrupted or missing databases with the copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run. After SQL is running again, you can restore any other databases, if needed.

If copies of the master and model databases were not made, then you must use Microsoft's rebuildm.exe utility to rebuild the master database and start SQL.

Because all changes made to the master database after the last backup was created are lost when the backup is restored, the changes must be reapplied. If any user databases were created after the master database was backed up, those databases cannot be accessed until the databases are restored from backups or reattached to SQL.

Note: It is recommended that you restore the master database first in a separate job. The SQL Server is restarted during the master database restore. If other database restores are included in the same job, the restores will fail.

To restore the SQL master database

- 1 On the **Backup and Restore** tab, in the list of servers view, right-click the server for which you want to restore the SQL master database, and then click **Restore**.
- 2 Select **Microsoft SQL Server databases**, and then click **Next**.
- 3 Select the backup set that contains the last master database backup, and then click **Next**.
- 4 Use the defaults or select the appropriate options and continue through the remaining wizard panels.
- 5 When prompted to run a consistency check, ensure that a check is run after the restore.

After the restore, SQL restarts in multi-user mode.

- 6 Restore the remaining SQL databases.

See [“Restarting SQL using database copies”](#) on page 1149.

Restarting SQL using database copies

You can restart SQL manually using copies of the database from previous backups and then restore the master database.

See [“Restoring the SQL master database”](#) on page 1148.

Table E-2 Restarting SQL using database copies

Step	Action
Step 1	Ensure that the SQL services are not running. Refer to the SQL Server documentation for details.
Step 2	Verify that the database copies are present. If necessary, restore the master and model database copies from a backup set to the same directory that the original master and model databases are in.

Table E-2 Restarting SQL using database copies (continued)

Step	Action
Step 3	Using the Windows Explorer, browse to the default data directory and delete the following files: ■ master.mdf ■ mastlog.ldf ■ model.mdf ■ modellog.ldf.
Step 4	Rename the copies of the databases back to their original names. Do not use read-only files. The SQL services will not start with read-only files.
Step 5	Use the SQL Service Control Manager to start SQL Server.
Step 6	Restore the latest changes to the master database.

The database copies are named master\$4idr, mastlog\$4idr, model\$4idr, and modellog\$4idr.

Table E-3 SQL database copy locations

SQL database copy	Location
An initial installation of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data*.*
A second installed instance of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.2\MSSQL\Data*.*
A default installation of SQL 2008	C:\Program Files\Microsoft SQL Server\MSSQL10.<instance name>\MSSQL\Data

The following table lists the copied database name and the original database name.

Table E-4 SQL database names

Copied database name	Original database name
master\$4idr	master.mdf
mastlog\$4idr	mastlog.ldf
model\$4idr	model.mdf
modellog\$4idr	modellog.ldf

Disaster recovery of a SQL Server

Backup Exec provides a quicker method for restoring SQL rather than running the Rebuild Master utility or reinstalling SQL to restart SQL. Using Backup Exec, you can replace the corrupted or missing databases with copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run.

If you use Simplified Disaster Recovery (SDR), then during an SDR recovery of drive C, it will automatically replace the damaged databases with the copies of the master and model databases. You can then restart SQL, and restore the latest master database backup and any other databases that are necessary.

The following topics are included in this section:

- See [the section called “How to prepare for disaster recovery of SQL”](#) on page 1151.
- See [the section called “Requirements for SQL disaster recovery”](#) on page 1152.
- See [the section called “Disaster recovery for an entire SQL Server or for SQL databases”](#) on page 1152.

How to prepare for disaster recovery of SQL

Do the following to prepare for disaster recovery of SQL:

- Back up both system and user databases and transaction logs regularly. Copies of the master and model databases are automatically created by Backup Exec whenever you back up the master and model databases. Backup Exec places these copies in the same directory that the databases are in, where they must remain in order to be updated.

The following table includes information about MS SQL database locations:

The copies of the master and model databases are named:

- Master\$4idr

- Mastlog\$4idr
- Model\$4idr
- Modellog\$4idr
- Back up the system drives that contain SQL instances.
Whenever you back up the system drive that contains a SQL instance, copies of the master and model databases are backed up. Backing up the system drive that SQL is on also backs up all the executables and registry settings needed for SQL to run.
- Back up the master database whenever any changes are made to SQL.
- Keep records of any service packs that have been installed.
- Ensure you are prepared to recover the entire server, not just SQL.

Requirements for SQL disaster recovery

To perform a recovery, you will need the following items:

- The latest backup of the SQL directory (\Program Files\Microsoft SQL Server\MSSQL), and the Windows registry/System State.
- The SQL database backups, and differential and log backups.
- An Administrator logon account (or an Administrator equivalent) during the recovery.

Disaster recovery for an entire SQL Server or for SQL databases

You can restore either the entire server, including the SQL databases, from full system backups, or restore only the SQL databases to a newly installed or other available SQL Server.

Restoring the entire server, including the SQL databases, has the added benefit of recovering other applications and data that may have resided on the server at the time of failure, and can be accomplished using one of the following methods:

- Manual recovery of the Windows server, and then manual recovery of the SQL databases. This method involves manually restoring the Windows server from full system backups, and then recovering the SQL databases.
- Simplified Disaster Recovery. This option provides an automated method of restoring the Windows server as well as the SQL databases from full system backups.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature”](#) on page 924.

To restore only the SQL databases, review the following:

- To restore only the SQL databases to a newly-installed or other available server, the server must be running on the same hardware platform (cross-platform restores are not supported), and the same version of SQL with the same service pack level as the original server.
- To restore SQL databases to an existing installation of SQL with other active databases, you should redirect the restore.
See [“Manual recovery of a SQL Server ”](#) on page 1153.

Manual recovery of a SQL Server

When you recover SQL manually, you must first restore the Windows server from full system backups. After recovery of the Windows computer is complete, or after the new server installation is available, you can recover the SQL databases.

For more information on how to run a manual disaster recovery, refer to the following sections:

See [“Performing manual disaster recovery of a local Backup Exec server on a Windows computer”](#) on page 930.

See [“Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer”](#) on page 935.

In order to restore SQL databases, SQL must be running; however, SQL cannot be started unless the master and model databases are present.

You can restore the master and model databases and start SQL using one of the following methods:

- Rename the files created by Backup Exec that replace the master and model databases. After the master and model databases are present on SQL, you must start SQL, and then restore all other databases.
See [“Restarting SQL using database copies”](#) on page 1149.
- Reinstall SQL.

This topic only details how to restart SQL by using the copies of the master and model databases made by Backup Exec. For more information on the Rebuild Master utility, or on reinstalling SQL, refer to your Microsoft SQL documentation.

If you are restoring to a new SQL installation, start with the restore of the master database.

See [“Restoring the SQL master database”](#) on page 1148.

About SQL Server Always On availability groups

Backup Exec supports SQL Server Always On availability groups. The Always On availability groups feature in SQL server is a high-availability and disaster-recovery solution. Backup Exec supports backup and restore of databases, which are part of Always On Availability groups.

The Always On availability groups feature maximizes the availability of a set of user databases by providing a fail over environment. The databases, known as availability databases, fail over together. An availability group supports a set of read-write primary databases and one to eight sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

Always On delivers capabilities such as, fail over cluster instances, multiple secondary, readable secondary, failover of group of databases (instead of entire node), and connection director. This feature can work in various customer environments with different requirements. An availability group fails over to an available replica.

Requirements of Always On availability groups from SQL Server

To set up Always on availability groups, there are many requirements and prerequisites from Microsoft SQL Server.

For more information about the requirements of Always On availability groups, refer to the Microsoft SQL Server documentation.

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/prereqs-restrictions-recommendations-always-on-availability?view=sql-server-ver15>

Terms used by SQL Server Always On availability groups

The following table describes the terms used by Microsoft SQL Server Always On availability groups.

Table E-5 Terms used by SQL Server Always On availability groups

Item	Description
Availability group	A container for a set of databases, availability databases, that fail over together.

Table E-5 Terms used by SQL Server Always On availability groups
(continued)

Item	Description
Availability database	A database that belongs to an availability group. For each availability database, the availability group maintains a single read-write copy (the primary database) and one to eight read-only copies (secondary databases).
Primary database	The read-write copy of an availability database.
Secondary database	A read-only copy of an availability database. Secondary databases are not backups. Continue to back up your databases and their transaction logs on a regular basis.
Availability replica	An instantiation of an availability group that is hosted by a specific instance of SQL Server and maintains a local copy of each availability database that belongs to the availability group. Two types of availability replicas exist: a single primary replica and one to eight secondary replicas.
Primary replica	The availability replica that makes the primary databases available for read-write connections from clients and, also, sends transaction log records for each primary database to every secondary replica.
Secondary replica	An availability replica that maintains a secondary copy of each availability database and serves as a potential fail over target for the availability group. Optionally, a secondary replica that can support read-only access to secondary databases supports creating backups on secondary databases.
Availability group listener	A server name to which clients can connect in order to access a database in a primary or secondary replica of an Always On availability group. Availability group listeners direct incoming connections to the primary replica or to a read-only secondary replica. This name is similar to cluster name and has associated IP address.

Table E-5

Terms used by SQL Server Always On availability groups
(continued)

Item	Description
Backup Preference	The Always On availability groups active secondary capabilities include support for taking backups on secondary replicas. Backup preference setting provides a way to specifies where backups should run. ■ Prefer Secondary Specifies that backups should occur on a secondary replica except when the primary replica is the only replica online. In that case, the backup should occur on the primary replica. This is the default option. ■ Secondary only Specifies that backups should never be performed on the primary replica. If the primary replica is the only replica online, the backup should not occur. ■ Primary Specifies that the backups should always occur on the primary replica. ■ Any Replica Specifies that backup jobs ignore the role of the availability replicas when choosing the replica to perform backups. Note: This setting is not enforced in Backup Exec.
Replica backup priorities	This setting allows giving a priority to each replica from 1 (minimum) to 100 (maximum) to select the node for the backup.
Exclude Replica	This setting allows to skip a node for performing backups. This is useful, for example, for a remote availability replica to which you never want backups to fail over.

Backup Exec recommendations for SQL Server Always On availability groups

Backup Exec recommendations to use SQL Server Always On availability groups for backups and restore.

- Set the backup preference as **Prefer Secondary**. You can then run the backup on the secondary node whenever possible and also run the job on the primary node in scenarios where, otherwise the job would fail. The user can then always have the latest data backed up.
- It is recommended that all the SQL Server instance names that are part of the availability group are the same, so that the database file path across all replicas is the same. It is also recommended that default data and log path on all nodes is the same. This ensures that the database file path on all nodes is always same.
If instance names are different, you can change the default data and log location to have the same paths for data and log files.
- It is recommended to automatically seed the initial data synchronization while adding a database in availability group. If the option is not available in all Microsoft SQL server versions, you can use the default option.
- It is important that database state is **Synchronized** on all nodes – not just the primary node. The backup may be successful, if the state is not **Synchronized**, but backup on secondary either fails or always runs on the primary node depending on the backup preference setting. For more information, refer to the SQL logs (**SSMS > Management > SQL Server Logs**) on the primary and secondary node.

Adding a listener to the list of servers on the Backup and Restore tab

You can add an availability group listener to the list of servers on the **Backup and Restore** tab so that you can select the SQL databases for backup.

To add a listener to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft SQL Server availability group**, and then click **Next**.
- 3 Select **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.

- 4 Enter the name of the availability group listener, and then click **Next**.

The name of the listener is available in SSMS (SQL Server Management Studio).

Note: Ensure that you can ping the FQDN and NetBIOS names of listener and all the participating nodes in the availability group from the media server.

Backup Exec displays the names of all the participating nodes of the availability group.

- 5 Click **Next**.

The Backup Exec Agent for Windows (RAWS) is installed on the participating nodes of the availability group. If the server already has RAWS installed, it is not installed again.

- 6 After RAWS is installed, click **Finish**.

The availability group listener is added in the servers list.

If a new node gets added later in availability group, you have to install RAWS on the new node. You can use same **Add server** wizard to install RAWS.

Backing up databases from a SQL Server availability group

Backup of SQL Server availability group databases is similar to backing up SQL databases and transaction logs. While backing up a SQL Server availability group, databases that are not part of the availability group are not displayed for selection.

See [“Backing up SQL databases and transaction logs”](#) on page 1127.

When you browse a node, all the databases available in the instance are displayed in the backup browse, but only those databases that are not part of the availability group are available for selection. Availability group databases are not available for selection when you browse any individual participating node for backup.

When you run a backup job, Backup Exec attempts to honour the backup preference setting, but there are some specific considerations from SQL.

- While checking the backup preference, Backup Exec also reads the priority of every replica. Backup Exec attempts to use the replica with the highest priority.
- Backup Exec checks if all replicas are in connected state and skips a replica that is not connected as per the availability group.
- Backup Exec checks and skips any replica that is marked for exclude from backup.

- Full and differential backups always run on the primary replica, irrespective of what is selected in the backup preference setting.

SQL server does not allow running Full and differential backups on secondary replicas.

A backup job may fail if you do not select the proper setting in backup preference. For example, if **Secondary Only** option is selected as the backup preference, and log incremental backup is running, but no secondary replica is in connected state, the job will fail.

If you select **Secondary Only**, some backup methods are not supported on secondary replicas. In such a scenario, Backup Exec can fail the job, but does not fail it where running backup on primary replica is the only option.

Information about the selected replica for backup is displayed in the backup job log.

It is recommended to run a full backup job after every restore, else the incremental is promoted to copy-only full job.

Restoring databases from a SQL Server availability group

Restore of SQL Server availability group databases is similar to restoring SQL databases and transactions logs. Restore always runs on the primary node.

See [“Restoring SQL databases and transaction logs”](#) on page 1146.

For SQL Server version 2014 and earlier, when the restore operation is complete, the database may be restored only on the primary node, and not added in the availability group. In such a scenario, the status on the database on the primary node is **Synchronizing** instead of **Synchronized**. On secondary replicas, the status is **Restoring**. SQL Server cannot get the database synchronized on all the nodes.

Perform the following steps in the order listed to get the database back in **Synchronized** state in all nodes:

1. Open SSMS (SQL Server Management Studio), connect to all the nodes, and browse to the availability group
2. In the primary node, select **Remove Database from availability group**.
3. Delete the database on all secondary replicas.

The database is only on the primary node, and it is not part of the availability group.

4. Right-click the availability group and click **Add database**.
5. Complete the wizard to add the database back in the availability group.

Backup Exec Agent for Microsoft Exchange Server

This appendix includes the following topics:

- [About the Backup Exec Exchange Agent](#)
- [Requirements for using the Exchange Agent](#)
- [Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations](#)
- [About installing the Exchange Agent](#)
- [Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab](#)
- [Managing preferred server configurations for Microsoft Exchange Database Availability Groups](#)
- [Recommended configurations for Exchange](#)
- [Requirements for accessing Exchange mailboxes](#)
- [Backup strategies for Exchange](#)
- [How Granular Recovery Technology works with the Exchange Information Store](#)
- [Snapshot and offhost backups with the Exchange Agent](#)
- [Backing up Exchange data](#)
- [Restoring Exchange data](#)
- [Disaster recovery of an Exchange Server](#)

About the Backup Exec Exchange Agent

The Backup Exec Agent for Microsoft Exchange Server (Exchange Agent) lets you integrate backups of Microsoft Exchange Server databases with network backups without separate administration or dedicated hardware.

The Exchange Agent provides the following features:

- Restore individual items from backups for which you enable Granular Recovery Technology.
- Restore to a PST file.
- Restore public folders.
- Restore a database to a specific drive and a path.
- Recreate a mailbox account.
- Search for a specific message to restore.
- Restore individual databases from snapshot backups by using the Recovery Database feature.
- Seeding of an Exchange Server database copy. Seeding adds a database copy to a location on another mailbox server in a database availability group (DAG).
- Off-host backup with Granular Recovery Technology (GRT) for Exchange Server.

For information about the best practices to use Backup Exec Agent for Microsoft Exchange Server, refer to *Backup Exec Best Practices*.

See [“About installing the Exchange Agent”](#) on page 1170.

See [“Backup strategies for Exchange”](#) on page 1176.

See [“Recommended configurations for Exchange ”](#) on page 1174.

See [“Granular Recovery Technology”](#) on page 739.

Requirements for using the Exchange Agent

The Backup Exec server must meet the following requirements:

Table F-1

Backup Exec server requirements for the Backup Exec Exchange Agent

Backup Exec server requirements	Description
To support the Exchange Agent	■ A license for the Backup Exec Agent for Microsoft Exchange Server (Exchange Agent) must be entered on the Backup Exec server. ■ The Backup Exec Agent for Windows must be installed on any remote Exchange Server that you want to back up. ■ The Backup Exec server must have access to the Exchange Server. You can find a list of compatible operating systems, platforms, applications, and supported service packs in the Backup Exec Software Compatibility List. ■ It is recommended that you use a Backup Exec services account that has local administrator rights on the Exchange server. You can have full Exchange permissions or minimal permissions that enable database backups and restores and Granular Recovery Technology operations. See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 1165.

Table F-1

Backup Exec server requirements for the Backup Exec Exchange Agent (continued)

Backup Exec server requirements	Description
To support Granular Recovery Technology (GRT) for the restore of individual items from Information Store backups	■ A Microsoft Windows operating system that supports minifilter drivers must be installed for Microsoft Exchange. You can find a list of compatible operating systems, platforms, applications, and supported service packs in the Backup Exec Software Compatibility List. ■ Storage that you use for GRT-enabled backups may have additional requirements. ■ If Exchange log files are on an advanced format disk or 512e disk, the Backup Exec server must also have a similar local volume available to perform GRT operations. See “Granular Recovery Technology” on page 739.

The following are requirements for the Exchange Server with the Backup Exec Exchange Agent:

Table F-2 Exchange Server requirements

Exchange Server requirements	Description
For operations on all Exchange Servers	The user account must be a member of the following group: ■ The Administrators group To support the Granular Recovery Technology feature, you must use the appropriate Exchange Server management utility to assign the user the required permissions. It is recommended that the user account have full Exchange permissions on the Exchange Server. If the user account cannot have full Exchange permissions, you can grant minimal permissions that enable database backups and restores and Granular Recovery Technology operations. See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 1165.

Table F-2 Exchange Server requirements (continued)

Exchange Server requirements	Description
To back up and restore Exchange Server	The following are requirements to back up and restore Exchange Server: ■ You must install the Agent for Windows on all of the mailbox server nodes in the DAG to back up the databases on a database availability group (DAG). ■ You must have local administrator rights on each node of a Database Availability Group and on the Microsoft Exchange mailbox server to back up and restore Microsoft Exchange database files. ■ You can configure minimal permissions for a user account to perform database backups and restores and to support Granular Recovery Technology on Exchange Servers if you cannot have full permissions. ■ For Microsoft Exchange Server 2013, you must have Microsoft Exchange Server 2013 with Cumulative Update 1 or later installed. See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 1165. See “About the Agent for Windows” on page 951.

For specific operating system requirements for the Exchange Agent, refer to the Backup Exec Software Compatibility List.

See [“Configuring Backup Exec to discover data to back up”](#) on page 713.

Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations

A user account must access mailboxes on the Exchange Server to perform Backup Exec operations. To gain access to the Exchange Server, you must have full access

permissions on the Exchange Server. If you want to restrict the access on the Exchange Server, you can grant minimal permissions that enable users to perform database backups and restores and Granular Recovery Technology (GRT) operations.

Ensure that the user account has local administrator's rights on the Exchange Server and then use one of the following methods to grant the permissions:

- Grant full access permissions at the Organization Administrators or Organization Management level. It is recommended that the user account have full Exchange permissions on the Exchange Server to perform Backup Exec operations.

Permissions are then propagated automatically to any new Exchange Servers that you add under the level at which the permissions are assigned.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

- If the user account cannot have full Exchange permissions for Backup Exec operations, you can grant minimal permissions. Minimal permissions let users perform database backups and restores and Granular Recovery Technology operations explicitly on each Exchange Server.

If you grant permissions explicitly and then add another Exchange Server, you must grant permissions explicitly on the added server as well.

[Granting minimal permissions for a user account to perform database backups and restores of Exchange Servers](#)

[Granting minimal permissions for a user account to support Granular Recovery Technology on Exchange Servers](#)

Granting minimal permissions for a user account to perform database backups and restores of Exchange Servers

You can grant minimal permissions for a user account that let you perform database backups and restores of an Exchange Server.

To grant full permissions for Microsoft Exchange, use an account with the Organization Management role.

To grant minimal permissions for a user account to perform database backups and restores of an Exchange Server

- ◆ Do one of the following:

To grant permissions for a user account using or the Exchange Admin Center in Microsoft Exchange

Add the user account to the following roles:

- Public Folder Management
- Recipient Management
- Server Management

To grant permissions for a user account using the Exchange Management Shell

Do the following in the order listed:

- **Type the following command:**

```
new-RoleGroup -Name <role
group name> -Roles
@("Database Copies",
"Databases", "Exchange
Servers", "Monitoring", "Mail
Recipient Creation", "Mail
Recipients", "Recipient
Policies" "Mail Enabled
Public Folders", "Public
Folders")
```

For example:

```
new-RoleGroup -Name
BackupExecRoles -Roles
@("Database Copies",
"Databases", "Exchange
Servers", "Monitoring", "Mail
Recipient Creation", "Mail
Recipients", "Recipient
Policies", "Mail Enabled
Public Folders", "Public
Folders")
```

- **Type the following command:**

```
Add-RoleGroupMember -Identity
<role group name> -Member
<name of the user account>
```

For example:

```
Add-RoleGroupMember -Identity
BackupExecRoles -Member
BackupExecUser
```

Granting minimal permissions for a user account to support Granular Recovery Technology on Exchange Servers

You can grant minimal permissions for a user account that let you support only Granular Recovery Technology (GRT) on an Exchange Server.

For more information about recipient scope, see the Microsoft Exchange documentation.

To grant permissions for a user account to support only Granular Recovery Technology on an Exchange Server using the Exchange Management Shell

1 Type the following command:

```
New-ManagementRole -Name "<management role name>" -Parent  
ApplicationImpersonation
```

For example:

```
New-ManagementRole -Name "EWSImpersonationRole" -Parent  
ApplicationImpersonation
```

2 Type the following command:

```
New-ManagementRoleAssignment -Role "<management role assignment  
name>" -User <user name> -Name "<assignment name>"
```

For example:

```
New-ManagementRoleAssignment -Role "EWSImpersonationRole" -User  
BackupExecUser -Name "BackupExecUser-EWSImpersonation"
```

3 Do the following:

For Exchange 2013 or later Type the following command:

```
New-ThrottlingPolicy -Name "<throttling policy name>" -EwsCutoffBalance Unlimited -EwsMaxBurst Unlimited -EwsMaxConcurrency Unlimited -ExchangeMaxCmdlets Unlimited -MessageRateLimit Unlimited -PowershellCutoffbalance Unlimited -PowershellMaxBurst Unlimited -PowershellMaxCmdlets Unlimited -PowershellMaxConcurrency Unlimited -PowershellMaxOperations Unlimited -RecipientRateLimit Unlimited -ThrottlingPolicyScope Regular
```

For example:

```
New-ThrottlingPolicy -Name "EWSRestoreThrottlingPolicy" -EwsCutoffBalance Unlimited -EwsMaxBurst Unlimited -EwsMaxConcurrency Unlimited -ExchangeMaxCmdlets Unlimited -MessageRateLimit Unlimited -PowershellCutoffbalance Unlimited -PowershellMaxBurst Unlimited -PowershellMaxCmdlets Unlimited -PowershellMaxConcurrency Unlimited -PowershellMaxOperations Unlimited -RecipientRateLimit Unlimited -ThrottlingPolicyScope Regular
```

4 Type the following command:

```
Set-Mailbox -Identity <user name> -ThrottlingPolicy "throttling  
policy name"
```

For example:

```
Set-Mailbox -Identity BackupExecUser -ThrottlingPolicy  
"EWSRestoreThrottlingPolicy"
```

5 Type the following command:

```
Set-ThrottlingPolicyAssociation -Identity <user name>  
-ThrottlingPolicy "throttling policy name"
```

For example:

```
Set-ThrottlingPolicyAssociation -Identity BackupExecUser  
-ThrottlingPolicy "EWSRestoreThrottlingPolicy"
```

About installing the Exchange Agent

The Exchange agent is installed as part of the Agent for Applications and Databases and can protect local or remote Exchange Server databases..

To support the Exchange Agent, the Backup Exec server must have access to the Exchange Server.

You can find a list of compatible operating systems, platforms, applications, and supported service packs in the Backup Exec Software Compatibility List.

Note: When you install Microsoft Exchange Tools and Backup Exec together on a server, Exchange Tools must be installed first. If you install Backup Exec before Exchange Tools, you must restart the Backup Exec server after you finish the Exchange Tools installation.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab

You can add an Exchange Server and a database availability group (DAG) to the list of servers on the **Backup and Restore** tab so that these servers can be selected for backup jobs. When you select the **Add** option in the **Servers and Virtual Hosts** group on the **Backup and Restore** tab, you can add Microsoft Windows computers and servers.

Note: When you add a Microsoft Exchange database availability group, it is recommended that you manually restart each Exchange Server after installing the Agent for Windows. If you select to automatically restart after install, all of your Exchange Servers in the DAG may restart at the same time.

To add an Exchange Server or DAG to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Do one of the following:

To add an Exchange Server

Click **Microsoft Windows computers and servers**.

To add a DAG

Click **Microsoft Exchange database availability group**.

- 3 Click **Next**.
- 4 Follow the **Add a Server** wizard prompts to add the Exchange Server or DAG to the list of servers in the **Backup and Restore** tab.

See [“Backing up Exchange data”](#) on page 1182.

Managing preferred server configurations for Microsoft Exchange Database Availability Groups

Preferred server configurations are collections of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where database copies are replicated

between multiple servers. You can create preferred server configurations for Microsoft Exchange Database Availability Groups (DAG).

You do not have to create a preferred server configuration to back up replicated database copies. You can let Backup Exec choose the best server from which to back up the replicated database copies. Designating a preferred server configuration gives you more control over your backup jobs. For example, you can select a local preferred server configuration to avoid having to back up replicated data over your WAN.

Backup Exec automatically includes the children of any site or DAG that you select as part of the preferred server configuration. To ensure that a backup is performed locally, you can select the local site as the preferred server configuration. Backup Exec selects from any of the local servers that belong to that site during the backup job. To ensure that a specific server is used for the backup, select only that server as the preferred server configuration.

You can create preferred server configurations for Microsoft Exchange Database Availability Groups. Preferred server configurations give you more control over your backup jobs since you can specify a preferred server from which Backup Exec backs up replicated data.

This topic contains information about the following subjects:

- [Creating preferred server configurations](#)
- [Deleting preferred server configurations](#)
- [Designating a default preferred server configuration](#)
- [Removing the default status for a preferred server configuration](#)

Creating preferred server configurations

To create preferred server configurations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Click **New**.
- 3 Select an Active Directory forest that contains the Exchange DAG members that you want to specify as preferred servers for backup.
- 4 Type a name for the group of preferred servers.
 You must enter a name before the preferred server configuration can be created.
- 5 Click **New** to enable the lists of available and selected servers from which you designate the preferred server.

- 6 In the **Available Servers** list, select the servers and sites that you want to use in the preferred server configuration.
- 7 On the **Preferred Servers** dialog box, click **OK**.
- 8 On the **Manage Preferred Servers** dialog box, click **OK**.

Deleting preferred server configurations

To delete preferred server configurations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration you want to delete.
- 3 Click **Delete**.
- 4 Click **OK**.

Designating a default preferred server configuration

You can designate a preferred server configuration to use as the default for all of your backup jobs that contain the appropriate replication data. When you back up data from a Microsoft Exchange Database Availability Group, you can set up Backup Exec to use your preferred server configuration that you selected as the default. You can override the selected preferred server configuration for specific jobs in the backup job settings.

When you designate a preferred server configuration to use as the default, it is not applied to existing backup jobs. The configuration is used for any subsequent backup jobs that you create.

To designate a default preferred server configuration

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration you want to set as the default.
- 3 Click **Set as Default**.
- 4 Click **OK**.

Removing the default status for a preferred server configuration

If you no longer want the preferred server configuration to be the default, you can remove its default status.

To remove the default status for a preferred server configuration

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration from which you want to remove the default status.
- 3 Click **Remove as Default**.
- 4 Click **OK**.

Recommended configurations for Exchange

Before starting backups for Exchange, read the following recommendations for configuring Exchange to make it easier to restore from backups:

Table F-3 Recommended configurations for Exchange

Recommendation	Description
Put transaction log files on a separate physical disk from the database.	This is the single most important configuration affecting the performance of Exchange. This configuration also has recovery implications, since transaction logs provide an additional recovery resource.
Make Write Cache unavailable on the SCSI controller.	The Windows operating system does not use buffers, so when Exchange receives a write complete notice from Windows, the write-to-disk has been completed. If Write Cache is enabled, Windows responds as though a write-to-disk has been completed, and will provide this information to Exchange (or other applications) incorrectly. The result could be data corruption if there is a system crash before the operation is actually written to disk.
Make circular logging unavailable if possible.	Circular logging minimizes the risk that the hard disk will be filled with transaction log files. But, if a solid backup strategy is in place, transaction log files are purged during the backup, thus freeing disk space. If circular logging is enabled, transaction log histories are overwritten, incremental and differential backups of databases are disabled, and recovery is only possible up to the point of the last full backup.
Avoid making the Exchange Server a domain controller.	For disaster recovery purposes, it is much easier to restore Exchange if you don't have to restore the Active Directory first.

Table F-3 Recommended configurations for Exchange *(continued)*

Recommendation	Description
Install Exchange into a domain that has at least two domain controllers.	Active Directory replication is not possible with only one domain controller in a domain. If the domain controller fails and corrupts the Active Directory, some transactions may not be recoverable if they were not included with the last backup. With at least two domain controllers in a domain, databases on the failed domain controller can be updated using replication to fill in missing transactions after the database backups have been restored.

See “[Requirements for accessing Exchange mailboxes](#) ” on page 1175.

Requirements for accessing Exchange mailboxes

Backup Exec must have access to a uniquely named mailbox within the Exchange organization for Information Store operations, depending on how the backup and restore jobs are configured.

Access to a uniquely named mailbox is required when you do the following:

- Configure a backup job that has all of the following settings:
 - A disk storage device other than a legacy backup-to-disk folder is the destination device.
 - The Granular Recovery Technology feature is enabled.
 - A backup method other than a snapshot method is used.
- You restore mailboxes and public folders.

You must use a Backup Exec logon account to connect to the Exchange Server when you select mailboxes or public folders for backup. Backup Exec attempts to find a mailbox with the same name as the user name that is stored in the Backup Exec logon account.

If you use a Backup Exec logon account that stores a unique user name and has a corresponding mailbox with the same name, then you are not prompted for an additional logon account. Otherwise, you must choose or create a Backup Exec logon account that stores the name of a unique mailbox within the Exchange organization.

A unique name does not share the first five characters in another mailbox name. For example, if EXCH1 is entered as the mailbox name, and there is another mailbox name such as EXCH1BACKUP, then Backup Exec cannot accept the name. You are prompted to choose another mailbox name.

You can choose or create a logon account that meets any of the following requirements:

- A logon account for which the user name matches a unique mailbox name.
- A logon account that uses a unique alias to a mailbox. The user account that connects to the Exchange Server must also have access to this mailbox.
- A logon account that uses the full computer name for a mailbox. The user account that connects to the Exchange Server must also have access to this mailbox.

An example of a full computer name is:

/O=Exchange_Organization/OU=Administrative_Group/CN=Recipients/CN=mailbox_name

See [“Granular Recovery Technology”](#) on page 739.

See [“Backup Exec logon accounts”](#) on page 759.

Backup strategies for Exchange

Backup Exec incorporates online, nondisruptive Exchange database protection as part of everyday backup routines, which increases the chance of data recovery and minimizes data loss without inhibiting daily activity. Backup Exec protects Exchange data, including the individual database, mailbox, and public folder with full, copy, incremental, and differential backups.

To decide which backup methods to use, consider the following:

- In small office environments with relatively small numbers of messages passing through the system, a daily full backup provides good data protection and the quickest recovery. If log file growth becomes an issue, consider using incremental online backups at midday to provide an added recovery point and manage the log file growth for you automatically.
- In large environments, incremental backups should be used to provide more frequent recovery point options throughout the day and to manage log file growth. Many organizations run full backups on a weekly basis, preferring to run incremental backups throughout the week to keep backup run time to a minimum. The trade-off with this technique occurs at recovery time when you must recover from the full backup and from each incremental backup as well.

What works best for you is based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

Consider the following backup strategies:

- Run full backups with the option to enable the restore of individual items selected so that you can restore individual mail messages and folders without restoring the entire database.

Depending on your environment, run full backups as follows:

- As frequently as possible, no less than once a day.
 - Daily with differential backups used at regular periods throughout the day.
 - Every few days (no less than weekly) with frequent incremental backups in between each full backup.
- Run Exchange backup jobs separately from other backup jobs.

In addition to backing up Exchange databases, you should also back up the following on a regular basis:

Table F-4 Backup selections for Exchange configuration data

Recommended backup selections for configuration data	Description
File system	Back up folders and drives containing files for Windows and Exchange. Usually, this is the root drive C:\, but may be different in each environment. Note: Back up the C:\ drive, but do not back up the virtual drive that is created by Exchange, if this virtual drive exists in your environment. It is intended only to provide Windows Explorer access to the Exchange data, but all file system functions may not be replicated. Backup and restore operations are not recommended or supported.
Windows registry	Back up the registry by running a full backup.

Table F-4 Backup selections for Exchange configuration data (continued)

Recommended backup selections for configuration data	Description
System State and/or Shadow Copy Components	Select System State and run a full backup to back up the following: ■ The Internet Information Service (IIS) metabase ■ The Windows registry See “About selecting data to back up” on page 171. If the entire server must be restored, you must restore both System State and Shadow Copy Components.
Active Directory	To back up Active Directory, select System State on the domain controllers and run a full backup. When there are configuration changes on the Exchange Server database, such as when objects are added, modified, or deleted, back up the Active Directory on the domain controllers. Note: Spread multiple domain controllers throughout each domain for efficient Active Directory replication, and so that if one domain controller fails, redundancy is still provided.

Note: Configure an Information Store backup for which the Granular Recovery Technology (GRT) feature is enabled to restore individual mailboxes, mail messages, and public folders.

See [“Granular Recovery Technology”](#) on page 739.

See [“Disaster recovery of an Exchange Server”](#) on page 1202.

How Granular Recovery Technology works with the Exchange Information Store

Backup Exec Granular Recovery Technology (GRT) lets you restore individual items from an Information Store backup without having to restore the whole backup. You should review the requirements for a GRT-enabled backup before you configure it.

See [“Granular Recovery Technology”](#) on page 739.

When you select items to restore from GRT-enabled backups, you cannot select the top level of the Information Store. To restore these items, you must restore the entire mailbox.

Note: When you perform a granular restore of a linked mailbox, shared mailbox, or site mailbox, do not select the **Recreate user accounts and mailboxes if they do not already exist on the destination server** option. These types of mailboxes must be created manually before you perform the restore. However, you can restore user accounts for linked mailboxes when you perform a granular restore of Active Directory from a backup of a domain controller. For more information about creating these types of mailboxes, see the Microsoft Exchange documentation.

You can also enable GRT when you create an off-host backup for the Information Store. Off-host backup lets Backup Exec move the backup process from the host computer to the Backup Exec server. The host computer is the remote computer that contains the volumes that you selected for backup. To run a GRT-enabled off-host backup, you must install the Backup Exec Advanced Disk-based feature on the Backup Exec server.

GRT and Microsoft Exchange Web Services

Backup Exec uses Microsoft Exchange Web Services (EWS) to support the Granular Recovery Technology feature. EWS provides support for the restore of individual mailboxes, mail messages, and public folders from an Exchange Server database backup.

Note: You must install .NET 4.0 and later, if not already installed.

To use EWS to restore individual items, Backup Exec disables the client throttling policy for the resource credentials you specify for the restore job. The client throttling policy enforces connection bandwidth limits on the Client Access server.

Backup Exec also creates an impersonation role and a role assignment for Exchange Impersonation. Exchange Impersonation role assignment associates the impersonation role with the Backup Exec resource credentials you specify for the restore job.

Backup Exec creates and assigns the following roles:

- EWSImpersonationRole
- EWSImpersonationRoleAssignment

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

See [“About the Advanced Disk-based Backup feature”](#) on page 1382.

Snapshot and offhost backups with the Exchange Agent

The Exchange Agent supports the Microsoft Volume Shadow Copy Service (VSS), a snapshot provider service. Using VSS, a point-in-time view of the Exchange database is snapped and then backed up, leaving the actual Exchange database open and available for users.

Offhost backup enables the backup operation to be processed on a Backup Exec server instead of on the Exchange Server. Moving the backup from the Exchange Server to a Backup Exec server enables better backup performance and frees the remote computer as well.

If the Advanced Disk-based Backup feature (ADBO) is installed on the Backup Exec server, you can use the Backup Exec Granular Recovery Technology (GRT) feature when you create an offhost backup for the Information Store.

See [“Setting default Granular Recovery Technology \(GRT\) options”](#) on page 745.

This topic includes information on the following subjects:

- [Configuring a snapshot backup for Exchange data](#)
- [Troubleshooting Exchange Agent snapshot and offhost jobs](#)

The Exchange Agent snapshot does not support the following:

- NAS configurations
- Mixing snapshot backups and non-snapshot backups

Due to a Microsoft Exchange limitation, if non-snapshot backups are run as part of a backup strategy, then snapshot backups should not be run. If snapshot backups are run, non-snapshot backups should not be done.

You can find a list of compatible storage in the Backup Exec Hardware Compatibility List.

The type of backup method that is available when using VSS with the Exchange Agent depends on the version of Exchange Server:

Table F-5 Available backup methods for Exchange snapshot versions

Exchange version	Available backup methods
Exchange Server	The following backup methods are available: ■ Full ■ Copy ■ Differential ■ Incremental snapshot backup ■ Individual database restore

Configuring a snapshot backup for Exchange data

You can run a snapshot backup for Exchange data.

Table F-6 Configuring a snapshot backup for Exchange data

Step	Action
Step 1	It is recommended that you perform consistency checks before running a snapshot backup. See “Backing up Exchange data” on page 1182.
Step 2	Create an Exchange backup job. See “Backing up Exchange data” on page 1182.
Step 3	If data that is not supported for snapshot backup is included in the backup selections, check Process logical volumes for backup one at a time to allow the job to complete with errors. This option is on the Advanced Disk-based Backup options on the Backup Job Defaults dialog box and on the Options dialog box for a backup job.

Table F-6 Configuring a snapshot backup for Exchange data *(continued)*

Step	Action
Step 4	Schedule or start the backup job. See “Backing up data” on page 158.

Troubleshooting Exchange Agent snapshot and offhost jobs

An Exchange Agent snapshot job fails on the following conditions:

- The Exchange Agent snapshot fails.
- If circular logging is enabled, and incremental or differential backup methods are selected.
- If an unsupported version of Exchange is installed. To verify supported versions of Exchange, refer to the Backup Exec Software Compatibility List.

Backing up Exchange data

To back up Exchange data, you can select the following:

- Individual databases
- Database availability groups (DAG)
You must enter an Exchange Agent license on the Backup Exec server for each Exchange Server in the DAG that you want to back up. You must then install the Agent for Windows on all the servers in the DAG.
Each database in the DAG must be backed up through the DAG container that displays in the list of servers on the **Backup and Restore** tab. The DAG container displays an Exchange logo on the server.

Note: If you add Exchange databases after you create a backup job, you must edit the backup job to include the new selections.

You can set backup job default options for all Exchange backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Automatic exclusion of Exchange data during volume-level backups

If you select a volume that contains Exchange data for backup, the Exchange Agent uses Active File Exclusion to automatically exclude Exchange data that should not be included in a volume-level backup. For example, .EDB and .STM

files, as well as transaction log files, should not be part of a volume-level backup because they are opened for exclusive use by Exchange.

Without this exclusion, these files appear as in use - skipped during a non-snapshot backup. During a snapshot backup, these files may be backed up in an inconsistent state, which could create restore issues.

While it is not recommended, if you want to include Exchange data in a volume-level backup, you must first dismount the databases that you want backed up, and then run the backup job.

To back up Exchange data

- 1 On the **Backup and Restore** tab, right-click the server that contains the Exchange data that you want to back up.

To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the resources that you want to back up and uncheck the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, click **Schedule**, and then select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, click **Storage**, and then select a storage device for this job.
- 9 On the **Backup Options** dialog box, in the left pane, select **Microsoft Exchange**.
- 10 Set any of the following options for this job:

Perform a consistency check before the backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider

Select this option to run a consistency check when the Microsoft Volume Shadow Copy Service option is selected. The option Microsoft Volume Shadow Copy Service is automatically used whenever a software backup is selected on the Advanced Disk-based Backup properties. You can also select the Microsoft Volume Shadow Copy Service on the Advanced Open File backup properties.

The consistency check, which is run on the snapshot, determines if possible data corruption exists.

If this option is selected, and the dependent option **Continue with backup if consistency check fails is not selected**, then data for specific Exchange objects that are determined to be corrupt are not backed up. All other non-corrupt Exchange objects are backed up.

For example, if a specific Exchange database file is corrupt, then backup is skipped only for that corrupt database file. All other non-corrupt database files and transaction log files are backed up.

When the option Continue with backup if consistency check fails is enabled, then all Exchange data is backed up regardless if corrupt files exist.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1180.

Continue with the backup if consistency check fails

Select this option to continue the backup job even if the consistency check fails. You may want the job to continue if you think a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database that may have only a small problem.

**Exchange in a Database
Availability Group**

Specify one of the following backup sources for Exchange:

■ **Back up from the passive copy and if not available, try the active copy (recommended)**

Lets you back up a passive copy of the database by default. Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings. However, if the passive copy is not available, Backup Exec backs up the active copy of the database. During the backup, database performance degradation can occur if you have to back up the database over a WAN.

■ **Back up from the passive copy only (job fails if not available)**

Lets you back up a passive copy of the database. If Backup Exec cannot access the passive copy, the job fails. In this case, neither the active nor the passive database is backed up. Select this option when you do not want to affect the performance of the active copy of the database.

Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings.

Note: You must have the preferred server settings configured to use this option.

■ **Back up from the active copy only (job fails if not available)**

Lets you back up the active copy of the database. If Backup Exec cannot access the active copy, the job fails. Therefore, neither the active copy nor the passive copy is backed up.

The active copy contains newer information than the passive copy. When you back up the active copy, you have a backup of the most recent database data.

■ **Back up from the preferred server configuration only (Use the passive copy first and if not available, use the active copy. Job fails if copies are not available in the preferred server configuration.)**

Lets you back up from the preferred server

configuration that you specify as the preferred backup source. Backup Exec selects the passive copy of the database as the backup source first. However, if the passive copy of the database is not available, Backup Exec selects the active copy as the backup source. If no copies of the database are available for backup on the configured preferred servers, the job fails.

See [“Managing preferred server configurations for Microsoft Exchange Database Availability Groups”](#) on page 1171.

Preferred server configuration Specify the preferred server configuration that you want to use for the **High Availability Server** options.

Click **Change** to create a new preferred server configuration or manage existing preferred server configurations.

See [“Managing preferred server configurations for Microsoft Exchange Database Availability Groups”](#) on page 1171.

Backup method

Specify one of the following backup methods that you want to use for this job:

■ **Full - Back up databases and logs (truncate logs)**

This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.

■ **Full Copy - Back up databases and logs**

This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.

You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups

Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.

Note: It is recommended that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Backup method

Specify one of the following Exchange-specific backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**

This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.

- **Full Copy - Back up databases and logs**

This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.

You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.

- **Differential - Back up logs**

This option backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after they are backed up.

To restore from differential backups, the last differential backup and the last full backup are required.

If circular logging is enabled, differential backups cannot be performed.

- **Incremental - Back up logs (truncate logs)**

This option backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database.

To restore from incremental backups, the last full backup and all the incremental backups done since the last full backup are required.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1180.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups

Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.

Note: It is recommended that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

11 Set any additional options for this job.

See [“Backing up data”](#) on page 158.

12 Click **OK**.

See [“Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab”](#) on page 1171.

See [“Changing default backup job settings”](#) on page 640.

See [“Editing backup definitions”](#) on page 205.

See [“Editing a stage”](#) on page 221.

Setting default backup options for Exchange Server

You can use the defaults that Backup Exec sets during installation for all Exchange Server jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Exchange Server

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for Exchange Server backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for the backup jobs that you send to different types of storage.

- 3 In the left pane, select **Microsoft Exchange**.
- 4 Select the appropriate options.

Perform a consistency check before the backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider

Select this option to perform a consistency check when the Microsoft Volume Shadow Copy Service option is selected. The option Microsoft Volume Shadow Copy Service is automatically used whenever a software backup is selected on the Advanced Disk-based Backup properties. You can also select the Microsoft Volume Shadow Copy Service on the Advanced Open File backup properties.

The consistency check, which is run on the snapshot, determines if possible data corruption exists.

If this option is selected, and the dependent option **Continue with backup if consistency check fails** is not selected, then data for specific Exchange objects that are determined to be corrupt are not backed up. All other non-corrupt Exchange objects are backed up.

For example, if a specific Exchange database file is corrupt, then backup is skipped only for that corrupt database file. All other non-corrupt database files and transaction log files are backed up.

When the option **Continue with backup if consistency check fails** is enabled, then all Exchange data is backed up regardless if corrupt files exist.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1180.

Continue with the backup if consistency check fails

Select this option to continue the backup job even if the consistency check fails. You may want the job to continue if you think a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database that may have only a small problem.

Exchange in a Database Availability Group

Specify one of the following backup sources for Exchange:

- **Back up from the passive copy and if not available, try the active copy (recommended)**

Lets you back up a passive copy of the database by default. Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings.

However, if the passive copy is not available, Backup Exec backs up the active copy of the database. During the backup, database performance degradation can occur if you have to back up the database over a WAN.

- **Back up from the passive copy only (job fails if not available)**

Lets you back up a passive copy of the database. If Backup Exec cannot access the passive copy, the job fails. In this case, neither the active nor the passive database is backed up. Select this option when you do not want to affect the performance of the active copy of the database.

Backup Exec selects the passive copy based on your selections in the Preferred Server settings.

Note: You must have the preferred server settings configured to use this option.

- **Back up from the active copy only (job fails if not available)**

Lets you back up the active copy of the database. If Backup Exec cannot access the active copy, the job fails. Therefore, neither the active copy nor the passive copy is backed up.

The active copy contains newer information than the passive copy. When you back up the active copy, you have a backup of the most recent database data.

- **Back up from the preferred server configuration only (Use the passive copy first and if not available, use the active copy. Job fails if copies are not available in the preferred server configuration.)**

Lets you back up from the preferred server configuration that you specify as the preferred backup source. Backup Exec selects the passive copy of the database as the backup source first. However, if the passive copy of the database is not available, Backup Exec selects the active copy as the backup source. If no copies of the database are available for backup on the configured preferred servers, the job fails.

See [“Managing preferred server configurations for](#)

[Microsoft Exchange Database Availability Groups](#) on page 1171.

Backup method

Specify one of the following backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups

Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.

Note: It is recommended that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

Backup method

Select one of the following Exchange-specific backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.
- **Differential - Back up logs**
This option backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after they are backed up.
To restore from differential backups, the last differential backup and the last full backup are required.
- **Incremental - Back up logs (truncate logs)**
This option backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database.
To restore from incremental backups, the last full backup and all the incremental backups done since the last full backup are required.
See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1180.

If circular logging is enabled, incremental and differential backups cannot be performed.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups

Select this option to enable the restore of individual items from incremental Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.

Note: It is recommended that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

5 Click **OK**.

Restoring Exchange data

The Exchange Agent lets you restore Exchange databases as well as individual mailbox items. You can restore items to their original location or you can redirect the restore to a new location.

Note: The **Restore Wizard** displays only up to 30,000 items. If you want to restore data from a database or mailbox that contains more than 30,000 items, you must search for the item that you want to restore. You can use search criteria such as the type of data and the date of the original backup to help reduce the number of items that display.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

The requirements and procedures for restoring Exchange data vary depending on the backup strategy you used. Before you restore Exchange data, you should review the required configuration and tasks.

You can restore Exchange data in the following ways:

- Use the recovery database to recover data from an older backup copy of the store without disturbing client access to current data.
- Restore Exchange data from snapshot backups.

- Restore individual Exchange items from a backup that is enabled for Granular Recovery Technology (GRT).
See [“How Granular Recovery Technology works with the Exchange Information Store ”](#) on page 1179.
- Restore Exchange data to a server other than the one from which it was backed up.

This topic includes the following information:

- [Requirements for restoring Exchange](#)
- [Configuring a destination database for restore data in Exchange](#)
- [Restore data by using the Exchange Server recovery database](#)
- [About restoring Exchange data from snapshot backups](#)
- [Restoring individual Exchange public folder messages from tape by duplicating backup sets](#)
- [About redirecting Exchange restore data](#)
- [About redirecting Exchange mailbox items](#)

Requirements for restoring Exchange

Review the following before restoring Exchange:

- The databases must already exist on the destination server, and must have the same names as the original databases.
- The destination server must have the same Organization and Administrative Group name as the source server.
- Before you start the restore, configure the destination databases so that they can be overwritten by a restore.
- You cannot restore messages with attachments that contain contact groups and undeliverable reports because Microsoft Exchange Web Services (EWS) for Microsoft Exchange 2013 does not support the restore of distribution lists and message reports when they are attachments. It is recommended that you restore messages that contain these types of attachments to a .PST file.

Configuring a destination database for restore data in Exchange

Before you restore Exchange, you should configure the destination database.

To configure a database

1 Do the following:

For Exchange 2013 or later Open the Exchange Admin Center.

2 Right-click the database that you want to overwrite.

3 Click **Properties**.

4 Do the following:

For Exchange Server On the **Maintenance** tab, select **This database can be overwritten by a restore**.

Restore data by using the Exchange Server recovery database

Exchange Server lets you mount a second copy of an Exchange mailbox store on any Exchange Server in the same Exchange Administrative Group as the original while the original store is still running and serving clients. This allows you to recover data from an older backup copy of the store without disturbing client access to current data.

Exchange Server uses recovery databases (RDB). Each server has a recovery database and there cannot be more than one mounted recovery database.

See your Microsoft Exchange documentation for more information about RDBs and recovery databases.

After the RDB is created, you can restore online backup sets to it. Then you can use the version of the Exchange Management Shell in Exchange to extract mailbox data from the stores into .PST files, and optionally merge the extracted data back into the online stores.

If the RDB resides on a different Exchange Server than the databases you are restoring, you should review the requirements for redirecting the restore of Exchange recovery databases.

[About redirecting Exchange restore data](#)

Following are requirements for restoring data using the Exchange Server recovery database:

- If multiple stores are selected for restore, mailbox stores in the RDB must come from the same database. You cannot add mailbox stores from different databases to the RDB at the same time.
- Public folder stores are not supported for restore using the RDB.

- Do not mount mailbox stores in the RDB before the restore. If you do mount the stores before the restore, then you must dismount them. Select the following option on the database property page in Exchange System Manager: This database can be overwritten by a restore.
Then, delete any files created in the data path for the RDB and added stores prior to restoring them.
Any files created in the data path for the RDB and added store or stores should be deleted as well, if stores were mounted prior to the restore.
- On the server that hosts the RDB, there must be a database with the same name as the original database for the data you are restoring. If no such database exists on the server, then you can use that name for the RDB when you create it.
- The Active Directory topology of the Exchange system must be intact and in the same state it was in when the backup was made. You cannot restore mailbox stores that were deleted and recreated. In addition, you cannot recover mailboxes from stores if the mailboxes were deleted and purged from the system or moved to other servers or mailbox stores.
- When the RDB exists on a server, the mailbox stores that it contains are the only stores that can be restored on that server by default. It is recommended that you create the RDB only when you intend to recover data using it, and remove the RDB from the server after the data recovery is complete.
- You can have more than one recovery database; however, you can only mount one recovery database to recover data.
- Do not mount the recovery database before the restore. If you do mount the recovery database before the restore, you must dismount it. Select the **This database can be overwritten by a restore** option on the database property page in the Exchange Management Console utility.

Refer to your Microsoft Exchange Server documentation for more information on the requirements and restrictions of recovering Exchange data.

About restoring Exchange data from snapshot backups

Note the following when restoring Exchange data from snapshot backups:

- If circular logging is enabled, only point-in-time, loss restores are possible. Roll-forward, no-loss restores cannot be performed.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 1180.

Restoring individual Exchange public folder messages from tape by duplicating backup sets

To restore individual public folder messages from tape, you must first duplicate the backup sets that contain the messages to disk storage. You can then restore the data from that disk storage.

The backup that you want to restore from must be a full backup or a copy backup. If there is an incremental or differential backup that is subsequent to the full backup, then you can restore individual items from the incremental or differential backup. The backup sets for the full backup and the incremental or differential backup must be on the same volume.

You cannot restore individual public folder messages from tape if the original backup is an incremental backup.

To restore individual Exchange public folder messages from tape by duplicating backup sets to disk

- 1 Insert the tape containing the required Exchange backup sets into a tape drive.
- 2 Create a duplicate backup sets job.
See [“Duplicating backup sets or a job history manually”](#) on page 222.
- 3 After the job completes, run a restore job to restore the individual public folder messages from the Exchange backup sets that are duplicated on the disk storage.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

About redirecting Exchange restore data

With Backup Exec, you can restore Exchange data to the server from which it was backed up or redirect the Exchange data to a different location. When redirecting Exchange data, the service pack on the Exchange Server where data is being redirected should be the same as the service pack on the original Exchange Server.

Following are requirements for redirecting Exchange database restores:

- The databases must already exist on the destination server.
- The destination server must have the same Organization Group name as the source server.
- The destination databases must be configured so that they can be overwritten.
See [the section called “Configuring a destination database for restore data in Exchange”](#) on page 1197.

You cannot redirect the restore of the following:

- A version of Exchange Server database to a different version of the database. Service packs for both Exchange Servers should also be the same.
- Site Replication Service (SRS) and Key Management Service (KMS). These services are dependent on the computer they reside on; redirection to another computer is not supported and could result in the loss of functionality of these services.

Note: KMS is not available in Exchange.

Before starting the redirected restore job, review information on finding and viewing specific data to restore, as well as for details on restore options and submitting restore jobs.

After completing the restore, it is recommended that a full backup of the restored databases be performed.

See [“Backing up Exchange data”](#) on page 1182.

About redirecting Exchange mailbox items

With Backup Exec, you can restore mailbox items such as mailboxes and public folders to a different mailbox on the same server or to a different location.

You can also restore mailboxes or mailbox items to a .PST, which is a Microsoft Outlook data file that is compatible with Microsoft Outlook (supported version). For supported Outlook versions, refer to the Backup Exec Software Compatibility List.

Following are requirements for redirecting Exchange mailbox and public folder items back to Exchange:

- The specified mailbox or public folder store must exist.
- The Backup Exec logon account must have rights to the destination mailbox.
- To overwrite existing public folder data, the logon account must have ownership of the public data. In this rare situation, it is recommended that when you restore public folder data, do not use the restore option **Overwrite existing mail messages and folders** on the Restore wizard panel titled **How do you want to restore the items?**

Following are requirements for restoring to a .PST file:

- The Agent for Windows must be installed on the destination server to which you redirect mailbox or public folder items.
- Microsoft Outlook (supported version) must be installed on the destination server.

Note: The maximum size of the .PST file is 20 GB. If the restore exceeds the size limit, the data is spanned across multiple .PST files which are numbered consecutively.

Disaster recovery of an Exchange Server

A disaster preparation plan is an absolute necessity for restoring Exchange efficiently and effectively in the event of a catastrophic failure. Because Exchange uses Windows security for authentication, disaster recovery of Exchange cannot be separated from the disaster recovery of Windows.

Planning ahead reduces the time needed to recover.

It is critical to build a kit that includes the following items:

- An operating system configuration sheet
- A hard drive partition configuration sheet
- Any RAID configuration
- A hardware configuration sheet
- EISA/MCA configuration disks
- An Exchange configuration sheet
- A Windows emergency repair CD

To perform the actual recovery, you will need the following items:

- An installed copy of Backup Exec
- The latest full, incremental, and differential backups of the Exchange databases you want to recover
- The Microsoft Exchange Server Installation CD
- Any service packs that were applied to the original installation

You can use Simplified Disaster Recovery to recover the Exchange Server.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CAS, Hyper-V hosts, and the Deduplication feature”](#) on page 924.

This procedure guides you through a complete restoration of Exchange using Backup Exec. You should have already performed all the appropriate preparation.

Always log on to Windows using the Administrator account (or an Administrator equivalent) during this procedure. Other requirements include:

- The databases must already exist on the destination server, and have the same names as the original databases.

- The destination server must have the same Organization and Administrative Group name as the source server.
- The destination databases must be configured so that they can be overwritten. See [“Restoring Exchange data”](#) on page 1196.

To perform disaster recovery for Exchange

- 1 Recover the Windows server first.
Ensure that you restore the Exchange Server files that existed on all disk partitions.
- 2 From the Services applet, verify the Microsoft Exchange Information Store service is started.
- 3 Start Backup Exec.
- 4 Catalog the backup sets of the Exchange Server databases you want to recover.
See [“Backup sets ”](#) on page 359.
- 5 Run the Restore Wizard and select the latest full backup set of each database for restore.
See [“Methods for restoring data in Backup Exec”](#) on page 235.
- 6 If necessary, select all subsequent incremental backup sets.
If differential backup sets are to be restored, only the most recent differential backup set needs to be selected.
- 7 After completing the restore, it is recommended that a full backup of the restored databases be performed.



Backup Exec Agent for Microsoft SharePoint

This appendix includes the following topics:

- [About the Agent for Microsoft SharePoint](#)
- [About installing the Agent for Microsoft SharePoint](#)
- [Requirements for the Agent for Microsoft SharePoint](#)
- [Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation](#)
- [Adding a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab](#)
- [Backing up Microsoft SharePoint data](#)
- [Restoring Microsoft SharePoint data](#)
- [Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec](#)
- [Viewing or changing SharePoint farm properties](#)
- [Disaster recovery of Microsoft SharePoint data](#)

About the Agent for Microsoft SharePoint

The SharePoint Agent enables network administrators to perform backup and restore operations on any supported Microsoft SharePoint installations that are connected to a network. SharePoint backups can be integrated with network backups without separate administration or dedicated hardware. The Agent for

Microsoft SharePoint is installed as part of the Agent for Applications and Databases.

For specific platforms that the SharePoint Agent supports, refer to the Backup Exec Software Compatibility List.

For information about the best practices to use Backup Exec Agent for Microsoft SharePoint (SharePoint Agent), refer to *Backup Exec Best Practices*.

See [“Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation”](#) on page 1206.

About installing the Agent for Microsoft SharePoint

Before you can back up Microsoft SharePoint server farms, you must install the Agent for Microsoft SharePoint (SharePoint Agent) on the Backup Exec server. The SharePoint Agent is installed as part of the Agent for Applications and Databases.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“Push-installing Backup Exec to remote computers”](#) on page 60.

See [“Requirements for the Agent for Microsoft SharePoint”](#) on page 1205.

Requirements for the Agent for Microsoft SharePoint

The Agent for Microsoft SharePoint (SharePoint Agent) has the following requirements:

- The SharePoint Agent must be installed on the Backup Exec server.
- The Agent for Windows must be installed on each remote SharePoint Server that you want to protect. In addition, the Agent for Windows must be installed on all servers in the server farm.
- You must use a logon account that has local administrative rights to back up and restore SharePoint data. The account should have local administrative rights on the servers where the SharePoint components are installed.
- The logon account that you use to restore content into an existing site collection must have appropriate rights to create objects in that site collection. If you restore into a site collection that does not exist, the logon account becomes the primary site collection owner.
- Internet Information Services (IIS) rights can affect database backups and restores. Ensure that the logon account that you use for backup and restore

has rights to access the IIS sites. Integrated Windows Security should be enabled within the IIS rights.

For specific operating system requirements for the SharePoint Agent, refer to the Backup ExecSoftware Compatibility List.

Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation

The Agent for Microsoft SharePoint includes support for Microsoft Office SharePoint Server, SharePoint Foundation, and SharePoint Server Subscription Edition.

SharePoint Server offers metadata features including tags, social bookmarks, and content ratings. These types of metadata are stored in the service applications that reside outside of the content database. For example, enterprise managed tags reside in the Managed Metadata Service application. You should back up all of your service applications to ensure that all metadata is protected.

Any metadata that is stored outside of the content database cannot be restored using Granular Recovery Technology (GRT). You can, however, use GRT to restore SharePoint data with metadata attached to it. As long as the metadata resides in the same service application, SharePoint maintains the link between the data and the metadata.

You can back up and restore the following types of SharePoint Server data:

- Web applications and their associated databases
- Individual documents and any pictures that are contained in libraries
- Sites and subsites
Individual objects and their versions can be restored from full database backups.
- Lists and list items
Individual objects and their versions can be restored from full database backups.
- Configuration database
A configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution when you restore this database. Any changes that you make to the farm topology before you restore from the backup are lost. The configuration database can be restored only to its original location.
- Service applications

See [“Disaster recovery of Microsoft SharePoint data”](#) on page 1218.

Adding a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab

Before you can back up Microsoft SharePoint data, you must add a SharePoint server farm to the list of servers on the **Backup and Restore** tab. If you add a single SharePoint server, Backup Exec adds the entire farm to which it belongs.

To add a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Do one of the following:

To add a single SharePoint server	Click Microsoft Windows computers and servers .
To add a SharePoint server farm	Click Microsoft SharePoint server farm .
- 3 Click **Next**.
- 4 Follow the **Add a Server** wizard prompts to add the SharePoint server farm or server to the list of servers in the **Backup and Restore** tab.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

Backing up Microsoft SharePoint data

The Agent for Microsoft SharePoint enables network administrators to perform backup operations on any Microsoft SharePoint installations that are connected to a network. SharePoint backups can be integrated with network backups without separate administration or dedicated hardware.

For more information about the specific types of SharePoint content that you can back up, see the following topics:

See [“Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation”](#) on page 1206.

Backup Exec's dynamic inclusion feature automatically protects any new resources that were added after a backup job was created. If Backup Exec discovers that you added a new resource as a child to a protected resource, it automatically

backs up the new resource. Because the backup job may include new resources, the job may require more storage space and more time to run than you anticipated.

You can set backup job default options for all SharePoint backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Note: Many backup options are available that are not described in the following procedure. For information on more backup options and methods, see the following topic:

See [“Backing up data”](#) on page 158.

To back up Microsoft SharePoint data

- 1 On the **Backup and Restore** tab, right-click the SharePoint server or farm that you want to back up.

To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the resources that you want to back up and uncheck the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, click **Schedule**, and then select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, click **Storage**, and then select the storage device that you want to use for the backup job.
- 9 On the **Backup Options** dialog box, in the left pane, click **Microsoft SharePoint**.
- 10 Set any of the following options for this job:

Perform a consistency check before the backup of any Microsoft SQL databases used by Microsoft SharePoint	Choose if you want to run a full consistency check (including indexes) of the Microsoft SQL databases that Microsoft SharePoint uses before you back up the databases.
Continue with the backup if consistency check fails	Choose to continue with the backup operation even if the consistency check fails.
Backup method	Specify one of the following backup methods that you want to use for this job: ■ Full - Back up databases Backs up the entire database. ■ Full Copy - Back up databases (copy) This option backs up the entire database without affecting future differential or log backups. Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup	Choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

Backup method

Specify one of the following SharePoint-specific backup methods that you want to use for this job:

- **Full - Back up databases**
Backs up the entire database.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
- **Differential - Back up database changes since the last full**
Backs up only the changes made to the database since the last full backup.
- **Differential (block-level) - Back up database changes since the last full**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full backup.
- **Incremental (block-level) - Back up database changes since the last full or incremental**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full or incremental backup.
- **Log - Back up and truncate transaction log**
Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated). If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup For this job, choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

See [“Granular Recovery Technology”](#) on page 739.

See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

- 11 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
 - 12 Click **OK**.
 - 13 In the **Backup Definition Properties** dialog box, click **OK**.
- See [“Changing default backup job settings”](#) on page 640.
- See [“Backing up data”](#) on page 158.
- See [“Editing backup definitions”](#) on page 205.

Setting default backup options for SharePoint

You can use the defaults that Backup Exec sets during installation for all SharePoint jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for SharePoint

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
 - 2 Select **Job Defaults**, and then select a backup option.
- For example, if you want to set up the default options for SharePoint backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for backup jobs that you send to different types of storage.
- 3 In the left pane, select **Microsoft SharePoint**.
 - 4 Select the appropriate options.

Perform a consistency check before the backup of any Microsoft SQL databases used by Microsoft SharePoint	Choose to run a full consistency check (including indexes) of the Microsoft SQL databases that Microsoft SharePoint uses before you back up the databases.
Continue with the backup if consistency check fails	Choose to continue with the backup operation even if the consistency check fails.
Backup method	Specify one of the following backup methods that you want to use for this job: ■ Full - Back up databases Backs up the entire database. ■ Full Copy - Back up databases (copy) This option backs up the entire database without affecting future differential or log backups. Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup	Choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

Backup method

Specify one of the following SharePoint-specific backup methods that you want to use for this job:

- **Full - Back up databases**
Backs up the entire database.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
- **Differential - Back up database changes since the last full**
Backs up only the changes made to the database since the last full backup.
- **Differential (block-level)**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full backup.
- **Incremental (block-level)**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full or incremental backup.
- **Log - Back up and truncate transaction log**
Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup For this job, choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job. See [“Granular Recovery Technology”](#) on page 739. See [“Configuring Instant GRT and full catalog options to improve backup performance for GRT-enabled jobs”](#) on page 662.

5 Click **OK**.

See [“Backing up Microsoft SharePoint data”](#) on page 1207.

Restoring Microsoft SharePoint data

The Agent for Microsoft SharePoint lets you restore the following types of Microsoft SharePoint data to the original location or you can redirect the restore to a new location:

- SharePoint individual items such as documents, images, sites, subsites, lists, and list items.
- SharePoint web applications or portal sites and their associated content.
- SharePoint farm components such as configuration databases, service applications, shared service providers, or other components.
- SharePoint individual items such as documents, sites, and list items located through search.

You should keep in mind the following things when you restore SharePoint data:

- When you restore SharePoint data, individual SharePoint documents are always restored to SharePoint document libraries as checked out. The documents are checked out using the same credentials as the logon account that you use for the restore. The documents must be checked in or published by that user before other users can use them.
- If you try to restore over a document that is checked out, the restore may fail. The restore fails if the document is checked out to a user that differs from the logon account credentials that are used for the restore.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

Restore individual items with Granular Recovery Technology (GRT)

Backup Exec also lets you restore individual documents, images, sites, subsites, lists, and list items from SharePoint database backups. To restore individual items from SharePoint database backups, ensure that the following Microsoft SharePoint option is selected during the backup job:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from the database backup

Note: Backup Exec does not support the restore of granular items from compressed databases or encrypted databases.

Redirected restores

Backup Exec lets you redirect the restore of the following:

- For SharePoint Server 2016 and later, Backup Exec allows restore of individual items such as documents and images on SharePoint Front End Web Server.
- Redirect SharePoint file-based data, such as documents and images that have been uploaded to a document library or are attached to list items, to an NTFS file system for SharePoint Server 2016 and later.

Backup Exec can extract and re-insert files directly into a SharePoint 2013 content database. However, the files cannot be selected for file system level redirection. SharePoint 2013 content databases files are stored as BLOBS, but in a shredded format. As a result, Backup Exec cannot redirect the restore of the shredded BLOBS to a file system.
- Redirect the restore of individual items such as documents, sites, and list items from one site to another site.

Note: This feature is not supported for SharePoint Server 2016 and SharePoint Server 2019.

When you redirect a restore from one site to another site, the restored items inherit the security permissions of the parent item to which they are restored. If the site is in a different SharePoint farm, the versions of SharePoint must be the same.

Note: You must use the same logon account for the original web server and the server in the site collection when you redirect the restore. You should also ensure that both of the servers use the same NetBIOS name, fully qualified domain name, or IP address.

- Restore SharePoint databases to an alternate SQL instance.
You can use the database to do the following:
 - Manually harvest data using the SharePoint Central Administration console in SharePoint.
 - Manually attach to a SharePoint web application.
- Redirect SharePoint web application content databases to an alternate web application.
The destination for the web application must be online and be of the same topology as the web application that was backed up. If the web application is in a different SharePoint farm, the versions of SharePoint must be the same.

Caution: When you restore SharePoint Portal document library data, any documents that exist in the selected destination and that have the same name as the documents being restored may be overwritten. You can select whether they should be overwritten in the properties for the restore job.

For more information about the different types of SharePoint content, see the Microsoft SharePoint documentation.

To restore Microsoft SharePoint data

- 1 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 2 Select **Microsoft SharePoint**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Using the Agent for Microsoft SharePoint with SharePoint Server and SharePoint Foundation”](#) on page 1206.

Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec

Backup Exec communicates with the web servers that participate in Microsoft SharePoint server farms to discover the farm topology. This process may take

some time if Backup Exec attempts to communicate with a web server that is unavailable. If you know that a particular web server in a farm is unavailable for a period of time, you can disable the communication between the web server and Backup Exec.

To disable or enable communication between a SharePoint web server and Backup Exec

- 1 On the **Backup and Restore** tab, double-click the SharePoint server farm to which the web server belongs.
- 2 In the left pane, click **Properties**.
- 3 Do one of the following:
 - To prevent Backup Exec from communicating with a SharePoint web server, clear the check box next to the web server name.
 - To let Backup Exec communicate with the SharePoint web server, select the check box next to the web server name.
- 4 Click **Apply**.

Viewing or changing SharePoint farm properties

You can view properties for any SharePoint farm that you monitor with Backup Exec, and you can change some properties. Backup Exec displays general and system information about the farm.

You can also enable or disable communication between a SharePoint web server and Backup Exec from the SharePoint farm properties dialog box.

See [“Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec”](#) on page 1216.

To view or change SharePoint farm properties

- 1 On the **Backup and Restore** tab, double-click the SharePoint farm whose properties you want to view.
- 2 In the left pane, click **Properties**.

3 View or change the appropriate properties:

SharePoint farm name	Displays the name of the SharePoint farm.
Description	Lets you enter a unique description to identify the farm in Backup Exec. The description is optional.
Logon account	Lists the logon account that Backup Exec uses to access the farm. Click Add/Edit to add a new logon account or to edit an existing logon account.
Web Servers	Lists the web servers that belong to the farm. You can enable or disable communication between the web servers and Backup Exec. See “Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec” on page 1216.

4 If you made changes, click **Apply**.

Disaster recovery of Microsoft SharePoint data

You can use the Agent for Microsoft SharePoint to recover a Microsoft SharePoint server after a hard drive failure. Before you recover SharePoint data, you must recover the SharePoint server's operating system.

You can use Backup Exec's Simplified Disaster Recovery feature or you can manually recover the server's operating system.

For more information about manual recovery, refer to the following sections:

See [“Performing manual disaster recovery of a local Backup Exec server on a Windows computer”](#) on page 930.

See [“Performing manual disaster recovery of a remote Backup Exec server or remote agent on a Windows computer”](#) on page 935.

See [“About Simplified Disaster Recovery”](#) on page 877.

After the Windows server is recovered, you can recover SharePoint data. Complete the actions in the table in successive order to recover the SharePoint data.

Table G-1 To recover SharePoint data after a disaster

Step	Action	Notes
Step 1	Recover the master database and the model database for the SQL instances that SharePoint uses. You must perform this step if you manually recovered the server's operating system.	Skip this step if you used Backup Exec's Simplified Disaster Recovery feature to recover the server's operating system. See "Manual recovery of a SQL Server" on page 1153.
Step 2	Inventory the media to be recovered.	See "Inventorying a storage device" on page 562.
Step 3	Catalog the media to be recovered.	See "Cataloging a storage device" on page 560.
Step 4	Restore the msdb databases for any SQL instances that SharePoint uses.	Select the backup sets that contain the msdb databases for the SQL instances that SharePoint uses. Configure the following Microsoft SQL restore options: ■ Use the default settings for all restore jobs. ■ Select Overwrite existing databases. See "Methods for restoring data in Backup Exec" on page 235.
Step 5	Restore all web applications.	Select the backup sets for all SharePoint web applications. Select Yes, restore over existing databases . See "Methods for restoring data in Backup Exec" on page 235.

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 6	Restore Shared Services Applications databases.	Restore the following Shared Services Applications databases: ■ Business Data Connectivity Service ■ Managed Metadata Service ■ PerformancePoint Service Application ■ Search Service Application ■ Secure Store Service ■ User Profile Service Application ■ Web Analytics Service Application ■ Word Automation Services ■ Services\State Services\Service DB 1 Select Yes, restore over existing databases. See “Methods for restoring data in Backup Exec” on page 235. Note: Some of the remaining restore jobs may fail because communication with the SharePoint server has not been fully established yet. This behavior is expected. Proceed with the recovery process until all steps are complete.

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 7	Restore search services.	Restore the following services: ■ SharePoint Foundation Help Search\Search Instance\Index Files 1 ■ Search-DB 1 Select Yes, restore over existing databases. See “Methods for restoring data in Backup Exec” on page 235. Note: You may get a message in the job log to restart your computer. You can ignore the message.
Step 8	Restore the SharePoint Configuration V4/V5-DB resource.	Select the backup sets for the SharePoint Configuration V4-DB resource. Select Yes, restore over existing databases. See “Methods for restoring data in Backup Exec” on page 235.
Step 9	Restart the SharePoint server.	After the restore job is complete, restart the SharePoint server. Then proceed to the next step.
Step 10	Restore remaining SharePoint resources.	Select the backup sets for the SharePoint Global Settings resources, if necessary. See “Methods for restoring data in Backup Exec” on page 235.

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 11 (For SharePoint multi-server farms)	Ensure that the servers in the SharePoint farm are configured correctly after the recovery.	Run the SharePoint Products Configuration Wizard.
Step 12	Back up the SharePoint server.	When the disaster recovery is complete, it is recommended that you perform a backup job as soon as possible. See “Backing up data” on page 158.

Backup Exec Agent for Oracle on Windows or Linux Servers

This appendix includes the following topics:

- [About the Backup Exec Oracle Agent](#)
- [About installing the Oracle Agent](#)
- [Configuring the Oracle Agent on Windows computers and Linux servers](#)
- [About authentication credentials on the Backup Exec server](#)
- [About Oracle instance information changes](#)
- [About backing up Oracle databases](#)
- [About restoring Oracle resources](#)
- [Best practices for Backup Exec Agent for Oracle on Windows and Linux Servers](#)

About the Backup Exec Oracle Agent

The Backup Exec Agent for Oracle on Windows or Linux Servers (Oracle Agent) uses Oracle's Recovery Manager (RMAN) to protect Oracle databases. RMAN is a tool that manages the backup and restore and recovery of Oracle databases.

The following features are available with the Oracle Agent:

- The ability to initiate backup and restore operations from Backup Exec or from the RMAN console as a Database Administrator (DBA).

Operations that the DBA performs on the RMAN console are referred to as DBA-initiated operations. You should refer to your Oracle documentation for information about RMAN.

- Multiple data stream support for increased performance during backup and restore operations.
- RMAN recovery catalog support to manage the backup, restore, and recovery of Oracle databases.
- Oracle Real Application Cluster (RAC) support.

For information about the best practices to use Backup Exec Agent for Oracle on Windows and Linux Servers (Oracle Agent), refer to *Backup Exec Best Practices*.

Oracle 12c and later supported version notes

Backup Exec supports following new features for Oracle 12c and later supported versions:

- Support for multitenant architecture
Backup Exec supports the new multitenant architecture introduced in Oracle 12c. In this new architecture, the Oracle database functions as a multitenant container database (CDB) that includes zero, one, or more pluggable databases (PDBs). A PDB is a user-created set of schemas, objects, and related structures that appears to an application as a separate database. All Oracle databases before Oracle database 12c were non-CDBs.
- New SYSBACKUP privilege for RMAN
For Oracle 12c and later supported versions, the Oracle Agent supports backup and restore tasks only with a user that has SYSBACKUP privileges.
- Support for a non-administrative user for ORACLE_HOME on Windows
For Oracle 12c and later supported versions, the Oracle services can run with a non-admin Oracle user. However, installation can be done only with the user that has administrator privileges.

The following are not supported:

- Tivoli Storage Manager (TSM) devices as storage for Oracle backup jobs.
- The Oracle Management Server.

See [“About installing the Oracle Agent”](#) on page 1225.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.

About installing the Oracle Agent

The Oracle Agent is installed as part of the Agent for Applications and Databases and protects local or remote Oracle instances.

To protect local or remote Oracle instances, you must install the following Backup Exec options:

- Backup Exec Agent for Windows on remote Windows computers.

Note: If you upgrade a previous version of the Agent for Windows on an Oracle server, you must restart the Oracle server after the upgrade. Backup Exec jobs cannot complete successfully until you restart the Oracle server.

See [“Methods for installing the Agent for Windows”](#) on page 68.

- Backup Exec Agent for Linux on remote Linux computers.
See [“About installing the Agent for Linux and Unix”](#) on page 1423.
- The Agent for Applications and Databases on the Backup Exec server.
See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

Configuring the Oracle Agent on Windows computers and Linux servers

Before you can back up or restore Oracle databases, you must do the following:

Table H-1 Configuring the Oracle Agent on Windows computers and Linux servers

Step	Action
Step 1	Configure information about the Oracle instances for the Oracle Agent. See “Configuring an Oracle instance on Windows computers” on page 1226. See “Configuring an Oracle instance on Linux servers” on page 1235.

Table H-1

Configuring the Oracle Agent on Windows computers and Linux servers *(continued)*

Step	Action
Step 2	Enable database access for the Backup Exec server. Whenever Oracle instance information changes or a new configuration is added, you must update the Backup Exec Agent Utility. If credential information is not updated, is incorrect, or the server is down, the error "Unable to attach to a resource..." may appear when you run a backup job. If this message appears, you must bring the server online and configure the information. For Oracle RAC, run the Backup Exec Agent Utility on each node and add information about the instances. When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Backup Exec Agent Utility. Note: When you use the Backup Exec Agent Utility, the user account with which you are logged on should be a member of the Oracle DBA group. You must have administrator privileges to run the Backup Exec Agent Utility. See "Enabling database access for Oracle operations on Windows computers" on page 1232. See "Enabling database access for Oracle operations on Linux servers" on page 1239.
Step 3	Set authentication credentials for Oracle. See "About authentication credentials on the Backup Exec server" on page 1240.

Configuring an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to configure information about the Oracle instances for the Oracle Agent on Windows computers.

To configure an Oracle instance on Windows computers

- 1
- On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.
- When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.
- 2
- On the **Oracle** tab, click **New**.
- Any instances that currently exist on the computer appear on the tab.
- 3
- Complete the appropriate options.
- See ["Oracle Agent Configuration options"](#) on page 1227.
- 4
- Click **OK**.

Oracle Agent Configuration options

You can set the following Oracle Agent Configuration options.

- See ["Configuring an Oracle instance on Windows computers"](#) on page 1226.
- See ["Editing an Oracle instance on Windows computers"](#) on page 1231.

Table H-2 Oracle Agent Configuration options

Item	Description
Local instance name	Displays the name of the Oracle instance. If you edit an instance, you cannot change the instance name. For Oracle RAC nodes, enter the fully qualified domain name of each physical node. The fully qualified domain name of the node appears in the list of servers on the Backup and Restore tab. The name is in the format RAC-<dbname>-<dbid>, where dbname is the database name, and dbid is the database ID.

Table H-2 Oracle Agent Configuration options (*continued*)

Item	Description
User name	Displays the user name for the Oracle instance. If the credentials for the Oracle instance change, you must enter a user with SYSDBA rights to the Oracle instance. For Oracle RAC nodes, enter the same set of credentials for all of the nodes. For an Oracle 12c and later supported database, you must enter a user name that has SYSBACKUP privileges.
Password	Displays the password for the Oracle instance user name.
Confirm password	Displays the password again to confirm it.
Oracle Home User will be granted the permissions to the Backup Exec Logs and Data folders	From Oracle 12c and later supported versions, the Oracle database supports the use of Oracle Home User. This user is a non-admin user; therefore, this user needs access rights to the Backup Exec Logs and Data folders. Note: This field is available only for Oracle 12c and later supported databases.

Table H-2 Oracle Agent Configuration options (*continued*)

Item	Description
Auxiliary instance path for PDB restores	For a point-in-time recovery of a pluggable database (PDB), RMAN first restores the PDB files from the appropriate backup. For the recovery of the PDB, RMAN needs a copy of the undo tablespace of the root container as it was in the specified point in time. To accomplish this recovery task, RMAN creates a temporary auxiliary database that consists of the root container's undo, system, and sysaux tablespaces. If the database that you want to restore uses a Fast Recovery Area, RMAN creates the auxiliary database files in this area in the <FRA> /<SID> /datafile directory. However, if this area does not have enough space, RMAN displays the following error: ORA-19809: limit exceeded for recovery files If the database that you want to restore does not use a Fast Recovery Area, RMAN creates the auxiliary database data files in the location you specify in the Auxiliary instance path for PDB restores field. The Agent Utility for Windows verifies the auxiliary path you specified. If this path is not available (non-existent), the utility creates an auxiliary path. If you do not specify this path, RMAN assumes that the Fast Recovery Area is configured and has sufficient space. Agent Utility cannot validate a path specified on the ASM disk and accepts the path as specified. Therefore, you must ensure that the path you entered is correct and accessible. Note: This field is enabled only for Oracle 12c and later supported databases.

Table H-2 Oracle Agent Configuration options (*continued*)

Item	Description
Use recovery catalog	Indicates that you plan to use the Oracle recovery catalog. The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the source database control file as the sole repository of metadata. The target for RMAN connection is either a target database (control file) or a recovery catalog. In Oracle 12c and later supported versions, RMAN connects to the container database if recovery catalog is not configured.
TNS name	Displays the Oracle Net Service name.
User name	Displays the user name for the Oracle recovery catalog.
Password	Displays the password for the Oracle recovery catalog.
Confirm password	Displays the password for the recovery catalog again to confirm it.
Backup Exec server name or IP address	Displays the name or IP address of the Backup Exec server where you want to send the DBA-initiated backup jobs. You must use the same form of name resolution for all operations.
Job template name	Displays the name of the Backup Exec job template that you want the DBA-initiated job to use for backup and restore operations. You create the job template on the DBA-initiated Job Settings dialog box on the Backup Exec server. If you do not specify a job template, the default job template is used. See “DBA-initiated job templates” on page 746.

Viewing an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to view information about the Oracle instances for the Oracle Agent on Windows servers.

To view an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.
- 2 On the **Oracle** tab, view the instances that currently exist on the computer.
See [“Oracle options for the Backup Exec Agent Utility”](#) on page 1231.
- 3 Click **OK**.

Oracle options for the Backup Exec Agent Utility

You can set the following Oracle options for the Backup Exec Agent Utility.

See [“Viewing an Oracle instance on Windows computers”](#) on page 1230.

Table H-3 Oracle options for the Backup Exec Agent Utility

Item	Description
Instance	Displays the name of the Oracle instance.
User Name	Displays the user name for the Oracle instance.
Recovery Catalog	Displays the name of the recovery catalog.
Backup Exec Server	Displays the name or IP address of the Backup Exec server where you want to send the DBA-initiated backup jobs.
Job Template	Displays the name of the DBA-initiated template. See “About performing a DBA-initiated backup job for Oracle” on page 1246.
New	Lets you add an Oracle instance.
Edit	Lets you revise an Oracle instance.
Delete	Lets you remove an Oracle instance.

Editing an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to revise information about the Oracle instances for the Oracle Agent on Windows computers.

To edit an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **Edit**.

Any instances that currently exist on the computer appear on the tab.

- 3 Edit the appropriate options.

See [“Oracle Agent Configuration options”](#) on page 1227.

- 4 Click **OK**.

Deleting an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to remove an Oracle instance for the Oracle Agent on Windows computers.

To delete an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **Delete**.

Any instances that currently exist on the computer appear on the tab.

- 3 Click **OK**.

Enabling database access for Oracle operations on Windows computers

You can use the Backup Exec Agent Utility to enable database access for the Windows computer after you configure an Oracle instance.

See [“About backing up Oracle databases”](#) on page 1243.

See [“About backing up Oracle RAC databases”](#) on page 1245.

See [“DBA-initiated job templates”](#) on page 746.

See [“Changing default backup job settings”](#) on page 640.

To enable database access for Oracle operations on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Arctera Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Database Access** tab.

(Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.

3 Complete the appropriate options to configure database access:

Enable the Backup Exec server to authenticate Oracle operations	Select this option to enable Oracle operations between the Backup Exec server and this computer.
User name	Specify a user name that has administrative rights to this computer. This logon account is what the Backup Exec server uses when it connects to this computer. If you specify an IP address or a fully qualified computer name as part of the user name, the Backup Exec Agent Utility may not be able to verify the user account. If the credentials entered are incorrect, the error “cannot attach to a resource” may be displayed when you run a backup or restore job. You must add this computer name and logon account to the Backup Exec server's list of authentication credentials for Oracle servers. If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you are browsing the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.
Password	Specify the password for this logon account. Note: For security reasons, the logon credentials are not stored on the remote computer.
Confirm Password	Type the password again to confirm it.
Use a custom port to connect to the Backup Exec server during Oracle operations	Select this option to change the port that is used for communications between this computer and the Backup Exec server during Oracle operations. By default, port 5633 is used. If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server.
Port number	Type the port number that you want to use for communications between this computer and the Backup Exec server.

4 Click **OK**.

- 5 For Oracle RAC installations, type the fully qualified domain name that you want to publish to.

The Backup Exec server that you publish to lists the RAC databases in the list of servers on the **Backup and Restore** tab.

If you do not enter a fully qualified domain name to publish to, the RAC databases are not in the list of servers.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.

- 6 On the Backup Exec server, add the name of the Oracle server and the user name that you enabled for database access to the Backup Exec server's list of authentication credentials.

See [“About authentication credentials on the Backup Exec server”](#) on page 1240.

Configuring an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to configure information about the Oracle instances for the Oracle Agent on Linux servers.

To configure an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.

Note: If the Oracle instance you want to protect is part of an Oracle RAC setup with version 12c or later supported version, switch to the Oracle user using the `su - <oracleuser>` command.

- 2 Change to the following directory:
cd /opt/VRTSralus/bin
- 3 Start the Backup Exec Agent Utility:
./AgentConfig
- 4 Type **2** to select Configure Oracle instance information, and then press **Enter**.
- 5 Type **1** to select the Add a new Oracle Instance option, and then press **Enter**.
- 6 Enter the name of the Oracle instance in upper case characters.
For example, ORACLENAME.

7 Enter the user name for the Oracle instance.

If the credentials for the Oracle instance are changed, you must update the credentials in this field. For Oracle RAC nodes, enter the same set of credentials for all of the nodes.

When you use the Backup Exec Agent Utility to enter the Oracle credentials for an instance, the credentials cannot be verified if the user account with which you are logged on is a member of the Oracle DBA group. If the credentials are incorrect, the error "Unable to attach to a resource..." may appear when you run a backup job.

Note: For Oracle 12c and later supported versions, the user should have SYSBACKUP privileges.

8 For Oracle 12c and later supported versions, when prompted, enter the auxiliary instance path for PDB restores.

If the database that you are restoring does not use a Fast Recovery Area, RMAN creates the auxiliary database data files in the location you specify.

Agent Utility cannot validate a path specified on the ASM disk and accepts the path as specified. Therefore, you must ensure that the path you entered is correct and accessible.

Note: The database instance is configured with the Oracle user; therefore, the Oracle user must have access rights on the directory specified in the auxiliary instance path. Else, the Agent Utility fails to create the auxiliary directory and you have to manually create the directory and assign read/write access to the Oracle user.

See ["Oracle Agent Configuration options"](#) on page 1227.

9 To display the Oracle database in a Backup Exec server's list of servers on the **Backup and Restore** tab, type the Backup Exec server name or IP address to which you want the remote computer to publish to.

- 10 When prompted, specify if you want to use a recovery catalog.

The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the source database control file as the sole repository of metadata.

The target for RMAN connection is either a target database (control file) or a recovery catalog. In Oracle 12c and later supported versions, RMAN connects to the container database if a recovery catalog is not configured.

If you specify a recovery catalog, any database that you want to back up must be registered in the recovery catalog before you can run backup jobs from the Backup Exec server.

- 11 To use a recovery catalog, type the recovery catalog name and a user name and password for the recovery catalog.
- 12 To use a customized DBA-initiated job settings template, type the name of the template.

See [“DBA-initiated job templates”](#) on page 746.
- 13 To commit the new entry to the configuration file, type **Y**, and then press **Enter**.

Viewing an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to view information about the Oracle instances for the Oracle Agent on Linux servers.

The following information is listed:

- Name of the instance
- Logon name for the instance
- IP address of the default Backup Exec server name for DBA-initiated operations
- Name of the DBA-initiated job template
- The auxiliary instance path for an Oracle 12c and later supported database (optional)

To view an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```

- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```

- 4 Type 4.

Editing an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to revise information about the Oracle instances for the Oracle Agent on Linux servers.

To edit an Oracle instance on Linux computers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.

- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```

- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```

- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**.
Any instances that currently exist on the computer are discovered.

Note: If the Oracle instance you want to protect is part of an Oracle RAC setup with version 12c or later supported version, then before selecting the **Configure Oracle instance information** option, switch to the Oracle user.

- 5 Type **2**.
- 6 Follow the prompts.

Deleting an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to remove an Oracle instance for the Oracle Agent on Linux servers.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.

To delete an Oracle instance for the Oracle Agent on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```
- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**.
Any instances that currently exist on the computer are discovered.
- 5 Type **3**.
- 6 Follow the prompts.

Enabling database access for Oracle operations on Linux servers

You can use the Backup Exec Agent Utility to enable database access for the Linux server after you configure an Oracle instance.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1241.

See [“About backing up Oracle databases”](#) on page 1243.

See [“About backing up Oracle RAC databases”](#) on page 1245.

See [“DBA-initiated job templates”](#) on page 746.

See [“Changing default backup job settings”](#) on page 640.

To enable database access for Oracle operations on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```
- 4 Type **1** to select Configure database access, and then press **Enter**.

- 5 Type the user name that is in the beoper group on the Linux system.

See [“About the Backup Exec operators \(beoper\) group for the Agent for Linux and Unix”](#) on page 1427.

If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you browse the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.

- 6 Type the password for this logon account, and then confirm it.

The logon credentials are not stored on this computer.

- 7 When prompted, specify if you want to use a custom port to connect to the Backup Exec server communications between this computer and the Backup Exec server during Oracle operations.

Port 5633 is used by default. If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server. If a Windows firewall is enabled, you must add this port as an exception.

See [“Changing network and security options for Backup Exec”](#) on page 718.

- 8 To commit the Oracle operation settings to the configuration file, type **Y**, and then press **Enter**.

About authentication credentials on the Backup Exec server

You must add the Oracle fully qualified domain name and the logon account name to the Backup Exec server's list of Oracle servers and authentication credentials. The Backup Exec server has database access for operations on Oracle instances that are included in the authentication list. Before you start any backup or restore operations, on the computer on which the Oracle instances are installed, ensure that you use the Backup Exec Agent Utility to configure instance information and database access.

The logon account name must have administrative rights to the Oracle server. If the user name is incorrect or is not provided, or if it does not have the appropriate rights, then you cannot perform Oracle backup or restore operations to that computer.

Note: For Oracle RAC nodes, enter the fully qualified domain name for the logon account name. You can view the fully qualified domain name of the node in the list of servers on the **Backup and Restore** tab. It is in the form RAC-<database name>-<database ID>.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1241.

See [“Deleting an Oracle server from the Backup Exec server’s list of authentication credentials”](#) on page 1242.

Setting authentication credentials on the Backup Exec server for Oracle operations

You must add the Oracle server to the list so that the Backup Exec server has database access for operations.

See [“About authentication credentials on the Backup Exec server”](#) on page 1240.

See [“About Oracle instance information changes”](#) on page 1243.

To set authentication credentials on the Backup Exec server for Oracle operations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Oracle**.
- 3 Enter the name of the Oracle server on which the instance is installed.

The name of the Oracle server should match the name of the server that lists the Oracle resource. It is recommended that you enter the fully qualified domain name. For example, Servername.domain.com is the fully qualified domain name and Servername is the NETBIOS name. For Oracle RAC nodes, enter the RAC-<database name>-<database ID> for each node in the list.

- 4 Click **Add**.

5 To add the logon account name, do one of the following:

Click the arrow Select the logon account name that you want to add.

Click <new logon account> Enter the appropriate options.

Use the same logon account format that you use when you enter the logon account name on the **Database Access** tab in the Backup Exec Agent Utility. For example, if you entered Domainname\Username on the Backup Exec Agent Utility, use that same format on the list of authentication credentials.

6 Click **OK**.

Oracle job settings options

You can add the Oracle fully qualified domain name and the logon account name to the Backup Exec server’s list of Oracle servers and authentication credentials.

See “[About authentication credentials on the Backup Exec server](#)” on page 1240.

Table H-4 Authentication credentials for Oracle servers options

Item	Description
Server name	Displays the name of the Oracle server.
Logon account	Displays the name of the logon account that has rights to the Oracle server.
Add	Lets you add the fully qualified domain name and logon account credentials to the list.
Delete	Lets you remove the fully qualified domain name and logon account credentials.

Deleting an Oracle server from the Backup Exec server’s list of authentication credentials

You can delete an Oracle server name or logon account from a Backup Exec server’s list of authentication credentials.

See “[About authentication credentials on the Backup Exec server](#)” on page 1240.

To delete an Oracle server from the Backup Exec server's list of authentication credentials

- 1 Click the Backup Exec button, and then select **Configuration and Settings** and then click **Backup Exec Settings**.
- 2 In the left pane, click **Oracle**.
- 3 Select the item that contains the server name or logon account that you want to delete.
- 4 Click **Delete**.
- 5 Click **OK**.

About Oracle instance information changes

Whenever information about the Oracle instance changes, such as the instance user name and password, you must update the Backup Exec Agent Utility.

When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Backup Exec Agent Utility. After these changes are entered, the Backup Exec server discovers them.

If the changes are not updated in the Backup Exec Agent Utility, the error "Unable to attach to a resource..." may appear when you run a backup job.

See ["Configuring the Oracle Agent on Windows computers and Linux servers"](#) on page 1225.

About backing up Oracle databases

Before you back up Oracle databases, review the following:

- You must run the Backup Exec Agent Utility on the Oracle server and add information about the instances before you can perform any backup or restore operations.
When Oracle instance information changes, you must update the Backup Exec Agent Utility. After these changes are entered, the Backup Exec server discovers them.
See ["Configuring the Oracle Agent on Windows computers and Linux servers"](#) on page 1225.
- During a backup operation, the amount of data that is backed up may not equal the size of the Oracle files that are on the disk. This behavior is normal. Backup Exec backs up the selected data files as well as a copy of the control file.

- In a Central Admin Server feature environment, all backup jobs for a specific Oracle instance must be delegated to the same managed Backup Exec server. If you do not restrict the backup job to the same managed Backup Exec server, then before you can restore data, you must move the physical media that contains the backup sets to a single managed Backup Exec server. See [“Selecting a Backup Exec server pool for backups”](#) on page 1364.

- If the Oracle database resides on the volumes that are configured with Oracle Automatic Storage Management (ASM), you cannot select these volumes as part of a file system backup.

The following message appears when you attempt to select the volumes:

```
An error was encountered while attempting to browse the
contents of <drive>. A device-specific error occurred.
```

- The database must be in a mounted or an open state before you can make backup selections.
- The database must be in ARCHIVELOG mode before the **Archived Logs** node can be displayed under the Oracle resource of the Oracle server.

Oracle 12c and later supported version notes

Before you back up an Oracle 12c and later supported database, review the following:

- To perform backup and recovery functions for an Oracle 12c and later supported database, the user must have the SYSBACKUP user privileges. For earlier database versions, the privilege and user for RMAN connection is SYSDBA and SYS.
- The target for RMAN connection is either a target database (control file) or a recovery catalog. In Oracle 12c and later supported versions, RMAN connects to the container database if a recovery catalog is not configured.
- Even if you select a single PDB or an individual tablespace or a data file for backup, the Oracle agent uses the CDB as the RMAN target if a recovery catalog is not configured. If you select multiple PDBs or an entire CDB or archived logs for backup, then the Oracle agent uses the CDB as the target. Thus, the target for RMAN is always the CDB. All the backups and restores happen by connecting to the CDB.
- If you select the **Oracle Database** node for backup, all data files of pluggable databases (including the data files of the root), archived logs, and control file are backed up.

If you select the **Pluggable Databases** node, the whole CDB is backed up including the archive logs. Archive logs are backed up only if the database is running in the Archive log mode.

- Each pluggable database node lists all the tablespaces of that pluggable database. If a PDB is selected for backup, all the tablespaces and the archived logs are also backed up along with the control file. Archive logs are backed up only if the CDB is running in the Archive log mode.

You can add an Oracle database to the list of servers on the **Backup and Restore** tab so that the database can be selected for backup jobs. You can set backup job default options for all Oracle backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for the particular job.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

See [“Backing up data”](#) on page 158.

See [“Oracle backup options”](#) on page 1247.

See [“About backing up Oracle RAC databases”](#) on page 1245.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1246.

About backing up Oracle RAC databases

Oracle Real Application Cluster (RAC) is an active and an active cluster with shared storage, in which multiple instances share a single physical database. Since all of the participating nodes can access the database, you can initiate backup, restore, or recovery from any node.

The Oracle RAC database resource container is added automatically to the Backup and Restore tab after

- You install the Remote Agent for Windows and configure the Backup Exec Agent Utility on a Windows Oracle server.
- You install the Remote Agent for Linux and UNIX and configure the AgentConfig utility on a Linux Oracle server.

Requirements for backing up Oracle RAC resources include the following items:

- You must run the Backup Exec Agent Utility on each node and add information about the instances before you can perform any backup or restore operations. When RAC nodes are added or removed, you must update the Backup Exec Agent Utility with information about the affected instances. After these changes are entered, the Backup Exec server discovers them.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.

- You must select the RAC fully qualified domain name when making backup selections.

Each node in the cluster uses the same fully qualified domain name. The fully qualified domain name of the node appears in the list of servers on the **Backup and Restore** tab. It is in the form RAC-**<database name>-<database ID>**.

Backing up Oracle RAC is similar to backing up standard Oracle databases.

You should be aware of the following differences:

- By default, each node in an Oracle RAC stores its archive logs locally. To have a meaningful backup of the archive logs, back up each archive log. Alternatively, you can move the archive logs to shared storage for backup.
- Each node that is part of the cluster is assigned a priority. For database backups, Backup Exec connects to the node that has the highest priority. Backup Exec uses the fully qualified domain name to connect to the node.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

See [“Backing up data”](#) on page 158.

See [“Oracle backup options”](#) on page 1247.

See [“About backing up Oracle databases”](#) on page 1243.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1246.

About performing a DBA-initiated backup job for Oracle

A Database Administrator (DBA) can initiate a backup or restore operation for Oracle from the RMAN console. Example scripts for backup and restore operations that you can run from the RMAN console are installed to the following location:

```
<Backup Exec install path>\Backup Exec\scripts\Oracle
```

Refer to your Oracle documentation for more information on using the RMAN console.

Review the following notes before initiating backup jobs for Oracle from the RMAN console:

- Ensure that you have completed all of the preparations for configuring the Oracle Agent.
See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 1225.
- The channel is not released if the RMAN console is not exited, or if a new manual channel is not allocated on that console.

- The SKIP INACCESSIBLE option is available in RMAN to skip corrupt data and log files. Jobs that include this option may complete successfully, but it is likely that if this data is restored, the database will be in an inoperable state. The SKIP INACCESSIBLE option is not available for Backup Exec server operations. If a backup job encounters corrupt data or log files, the job fails. It is recommended that you do not use this option.
- In a Central Admin Server feature (CAS) environment, the destination storage that you select in the DBA-initiated job template must be locally attached to the central administration server.
If the destination storage includes a storage device pool, all storage in the pool must be locally attached to the central administration server.

See [“About Oracle instance information changes”](#) on page 1243.

See [“About Oracle instance information changes”](#) on page 1243.

See [“About backing up Oracle databases”](#) on page 1243.

See [“DBA-initiated job templates”](#) on page 746.

Oracle backup options

The following options are available for Oracle backup jobs. These options appear when you select the Oracle option on the **Backup Job Defaults** dialog box and on the **Backup Options** dialog box for a backup job.

See [“About backing up Oracle databases”](#) on page 1243.

See [“Backing up data”](#) on page 158.

See [“Changing default backup job settings”](#) on page 640.

Table H-5 Oracle backup options

Item	Description
Maximum number of devices to use for resources that support multiple data streams	Specifies the maximum number of devices that the backup job can use. If you specify more than one device, you must choose one of the following items as the destination storage for the backup job: ■ A storage pool. ■ A legacy backup-to-disk folder that has at least two concurrent operations enabled. If there is only one storage device for the backup job to use, then the data streams from RMAN are backed up serially to the media. This feature is not available for DBA-initiated jobs.
Minimum number of devices that is required, fail the job if fewer devices are available	Specifies the minimum number of storage devices that the job can use. If the job cannot acquire the minimum number of devices, the job fails. This feature is not available for DBA-initiated jobs.
Delete archived log files	Lets you delete the archived log files automatically after the backup.
Do not back up archived log files that have already been backed up	Enables Backup Exec to skip any archived log files that have been backed up previously.
Perform the backup offline	Enables Backup Exec to take the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete. Select this option if the Oracle database is a non-archived logged database.

About restoring Oracle resources

The Oracle Agent lets you restore Oracle databases, tablespaces, or datafiles. You can restore items to their original location or you can redirect the restore to a new location. The restore selections that you choose in Backup Exec are converted to a script. RMAN uses the script to determine what to restore from the Backup Exec backup set. After the data has been restored to the Oracle server, RMAN

completes any requested recovery and restore operations. The options you select determine the recovery and restore operations.

Some recovery operations may not require media from the Backup Exec server. For example, the redo logs may still be on the Oracle server. During a restore operation, the amount of data that is restored may not be equal to the amount of data that is backed up. In some cases, the amount of data that is restored is listed as 0 bytes. This behavior is normal because Oracle might skip the datafiles that are already up-to-date on the disk.

Note: Backup Exec Agent for Oracle does not support the restore of a resource if you back up the resource on a different storage media. For example, you cannot restore an Oracle database if you have done a full backup of the Oracle database using the Backup-to-disk option on a device that is attached to the Backup Exec server and an incremental backup of an Oracle database on a Remote Media Agent device.

If you perform a complete recovery on the whole database, or on a tablespace or datafile, you must restore a backup of the database or files that you want to recover. Then you must apply online or archived redo logs, or both. For the jobs that are initiated both from the Backup Exec server and from a DBA, RMAN determines the specific data that it requires from Backup Exec to complete the restore and the recovery that you request.

Note: Backup Exec does not support Oracle tablespace point-in-time restore (TSPITR) through server-initiated operations.

Backup Exec does not support restore from storage device media and from backup sets for the Oracle Agent because during restore, RMAN determines which backup sets are needed for the job.

For Oracle 12c and later supported databases, Backup Exec supports Oracle pluggable database point-in-time restore (DBPITR) through server-initiated and DBA-initiated operations.

The point-in-time (PIT) restore job of the root fails with an error message: "Specifying CDB\$ROOT database is not supported". This feature is not supported by Oracle.

For Oracle 12c and later supported versions, root is displayed as a **Tablespaces** node in the backup browse and restore view. To restore the root, select the **Tablespaces** node under the control file node in the **Details File** view.

You can recover only the root if a data corruption or user error occurs that affects only the root. However, Oracle recommends that you recover all PDBs after

recovering the root to prevent metadata inconsistencies among the root and the PDBs. In such a case, it is preferable to perform a complete recovery of the whole CDB.

The Oracle agent retains the state of the pluggable databases (PDBs) as they were before the restore job. After the restore, the agent brings the PDBs to the same state as they were before the restore job.

Note: All states such as OPEN, MOUNTED, READ-ONLY, and READ/WRITE are retained except MIGRATE. If any PDB was in the MIGRATE state before the restore job, then after the restore that state of the PDB changes to MOUNTED.

You can only choose Oracle restore selections from the **Resource View** in the Restore Wizard. The **Details View** displays backup sets, but you cannot browse or select the contents.

On the **Resource View**, you can make restore selections from the online database or from control files.

Table H-6 Restore selections for Oracle resources

View restore data in	Description
Online database	Provides a view of the live database (if available). You can select an entire database or select individual tablespaces and datafiles. For an Oracle 12c and later supported database, this view displays the pluggable databases, the Tablespaces node, and the archived logs at the same level. In this view, you can select the whole tablespace and individual data files for restore, but you cannot select the pluggable databases and the archived logs nodes for restore. For Oracle 12c and later supported versions, Backup Exec supports restore and recovery of one or more tablespaces and data files of a PDB. Redirect restore of a PDB on the same host to a different path is supported. However, point-in-time restore of tablespaces and redirected restore of a PDB to a new host are not supported. Note: For Oracle RAC, the Oracle database is listed under its fully qualified domain name. It is in the form RAC-<i>database name</i>-<database ID>.

Table H-6 Restore selections for Oracle resources (*continued*)

View restore data in	Description
Control files	Provides a list of all backed up control files. Each control file lists the date it was backed up and the control file's piece ID. You cannot select individual tablespaces or datafiles for restore. For an Oracle 12 database, in this view, you can select only the control file node and the pluggable database nodes for restore. If you select the control file node for restore, the entire CDB is restored. This restore data includes all data that is related to the CDB and all PDBs. This data is restored and recovered to the PIT specified in the restore job. This feature is the same as the control file restore of a non-CDB database. If you select an individual pluggable database, only the selected PDB is restored. Note: When you recover to a point in time by using a control file, ensure that the date of the control file backup is before the specified recovery point in time. There should not be any database structure changes between the two times. Additionally, when you restore a control file, the entire database reverts to the point in time of the restored control file.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

About DBA-initiated restore for Oracle

DBAs can initiate restore jobs directly from the RMAN console. For example, you can specify the resources you want restored, and the number of channels to allocate for the restore job. Refer to your Oracle documentation for more information on using the RMAN console.

All DBA-initiated restore jobs are deleted after the jobs have completed.

Note: If you attempt to use a DBA-initiated restore job to restore a datafile, a tablespace, or a database that is online, a message appears on the RMAN console. The message indicates that the restore cannot be performed because Oracle does not allow the restore of these items if they are online. However, this message is not reported to Backup Exec. Therefore, the DBA-initiated restore job is reported in Backup Exec as completing successfully.

Oracle restore options

The following options are available when you restore Oracle data.

Table H-7 Oracle restore options

Item	Description
To the latest available time	Restores the Oracle database to the most recent full and incremental backups that are available.
To a point in time up to and including the specified time	Restores data up to and including a point in time. After the point in time, recovery stops. Enter a date and time or click the arrow to display a calendar from which you can select a date and time.
To an SCN	Restores up to and including a specific system control number (SCN). Type the SCN in the field provided.

About redirecting a restore of Oracle data

In Backup Exec, you can redirect an Oracle instance or its files by redirecting the following:

- An Oracle instance to another Oracle server.

Note: If you redirect the instance to a different Oracle server, ensure that an instance with the same name and database ID (DBID) is set up on that server. The database status should be Nomount. Refer to your Oracle documentation for details on creating an instance with the same name and database ID.

- An Oracle instance to another Oracle server and specifying alternate paths for the Oracle files.
- Tablespace, datafiles, and archive logs to an alternate location on the original server.
- For Oracle 12c and later supported versions, Backup Exec supports redirected restore of a CDB and redirected restore of a PDB on the same host to a different path. However, Backup Exec does not support redirected restore of a PDB to a new host.

It is recommended that you select only one instance for each redirected restore operation.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

Oracle advanced restore options

The following advanced options are available when you restore Oracle data.

Table H-8 Oracle advanced restore options

Item	Description
Open the database after the restore	Ensures that the database is opened as soon as the recovery is finished. Check this option if you want the database to be online after the recovery.
Restore read-only files if they are not current	Enables RMAN to examine the headers of all read-only data files and restore any that are not current.
Validate only; do not restore data	Mounts all required media and reads it as necessary. RMAN selects the backup sets that are necessary to perform the operation, and scans them all to ensure that they are available and not corrupted. No data is written or restored to the database server. Validation of the control file is not supported. It is recommended that you select this option to ensure that all required media is available before you attempt to restore to the database.
Restore data if validation completed without errors	Lets you run the restore job immediately if the validation was successful. All options that you have selected for restore and recovery are performed.
Restore only the control file	Recovers the control file for the Oracle database, but does not include the tablespaces or the associated data files.
After the restore, delete the archived logs that are no longer needed	Deletes older archived redo log files and free space on the hard disk.

Table H-8 Oracle advanced restore options (*continued*)

Item	Description
Maximum number of devices to use for resources that support multiple data streams	Specifies the maximum number of devices that the restore job can use. If you specify more than one device, you must choose one of the following items as the destination storage for the backup job: ■ A storage pool. ■ A legacy backup-to-disk folder that has at least two concurrent operations enabled. If there is only one storage device for the backup job to use, then the data streams from RMAN are backed up serially to the media. This feature is not available for DBA-initiated jobs.

Performing a redirect restore of an Oracle 12c and later supported database using Backup Exec

Redirect restore database scenarios

- Database using an OMF (Oracle Managed File)
 - Restoring the database to the same path on a new host
Before running the restore job, ensure that the folder structure on the new Oracle server is same as the original database. In case of an OMF database, PDB folders are created as GUIDs; therefore, you have to manually create the same folder structure as the original database.
 - Restoring the database to a different path on a new host
You can specify the redirect restore path in the Backup Exec UI. Oracle Agent creates the folder structure using PDB names in the redirect restore path you specified and restores the data in the appropriate folders. For example, the redirect restore path you specified in Backup Exec is `C:\Redirect`. The CDB that you want to restore has two PDBs namely, PDB1 and PDB2. These PDBs have GUID-based folders in the original database. After the restore, data is restored to the following locations:
 - PDB1 data is restored in `C:\Redirect\PDB1`
 - PDB2 data is restored in `C:\Redirect\PDB2`
 - PDB\$SEED data is restored in `C:\Redirect\PDB$SEED`
 - ROOT data is restored in `C:\Redirect\ROOT`

Note: If the Oracle agent fails to create the directory structure, the restore job fails with an error asking the user to create the structure manually.

- Database not using an OMF
 - Restoring the database to the same path on a new host
Before running the restore job, ensure that the folder structure on the new Oracle server is same as the original database.
 - Restoring the database to a different path on a new host
You can specify the redirect restore path in the Backup Exec UI. Oracle Agent creates the folder structure using the PDB names in the redirect restore path you specified and restores the data in the appropriate folders. For example, the redirect restore path you specified in Backup Exec is `C:\Redirect`. The CDB that you want to restore has two PDBs namely, PDB1 and PDB2. These PDBs have name-based folders in the original database. After restore data is restored to the following locations:
 - PDB1 data is restored in `C:\Redirect\PDB1`
 - PDB2 data is restored in `C:\Redirect\PDB2`
 - PDB\$SEED data is restored in `C:\Redirect\PDB$SEED`
 - ROOT data is restored in `C:\Redirect\ROOT`

Note: If the Oracle agent fails to create the directory structure, the restore job fails with an error asking the user to create the structure manually.

Note: Backup Exec does not support redirect restore of a PDB to a new host. However, Backup Exec supports redirect restore of a PDB on the same host to a different path.

Points to remember in case of a redirect restore

- By default, RMAN restores the datafiles in the same location from which it was backed up.
- The database directories and paths are not created from RMAN, but they must pre-exist.
- Delete the redo logs before starting the restore job because on-disk redo log are not useful. If you do not delete the redo logs, restore job might fail with the following error:

ORA-19698: C:\APP\ORAUSER\ORADATA\STRING\REDO01.LOG is from different database: id=nnnnnn, db_name=STRING

- If you want to restore the database to the same path on a new host, make sure that the directory structure is same as the original database. However, if the redirect restore is to a different path on a new host, the Oracle Agent creates the directories on the new path.

Follow the procedure to use a new Oracle server to recover the complete Oracle instance or database after it has been lost, deleted, or destroyed.

Performing a redirect restore of an Oracle 12c and later supported database using Backup Exec

- 1 Recreate the Oracle database using the same name of the original database that is no longer available.
- 2 Locate and rename the `pwd<SID>.ora` file. Make note of the path to this file for use in Step 3b.
- 3 Create a new `pwd<SID>.ora` file using the following steps:

- Open a command prompt.

- Type the following command

```
orapwd file="path from Step 2"\pwd<sid>.ora password=<password>
```

For example, C:\> orapwd file=c:\oracle\product\12.1.0\db_1
\database\pwdORCL1.ora password=NEWpassword

Note: Backup Exec Remote Agent must be installed and configured for Oracle database access.

- 4 In the command prompt, type the following command

```
SQLplus / as SYSDBA
```

- 5 Type `SHUTDOWN IMMEDIATE;`

- 6 Type `STARTUP NOMOUNT;`

- 7 Exit `Sqlplus` and launch `RMAN` from the command prompt.

- 8 Type `SET DBID=dbid ID;`

Note: This `dbid` must be identical to the `dbid` from the original Oracle instance. When complete, type `EXIT`.

- 9 On the Backup Exec server, on the navigation bar, click the arrow next to **Restore**, and then click **New Restore Job**.
- 10 In the **Properties** pane, under **Source**, click **Selections**.
- 11 Select the appropriate control file to restore.
- 12 In the **Restore job properties** pane, under **Destination**, click **Oracle Redirection**.
- 13 Select the **Restore Oracle instance to server** check box. Enter account credentials to access the new or the alternate Oracle server.

If the Oracle datafiles and archive logs are stored on a different location on the target server, enable the **Redirect Oracle files to path** option, and specify the valid paths for the datafiles and archive logs. Click **Run Now** to run the restore job.

Note: The restore job fails because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.

- 14 On the Oracle server, start SQLplus / as SYSDBA from the command prompt.
- 15 Type `alter database open resetlogs;`

If an error is encountered while Oracle tries to open the database, document the online redo log path, and then update the path using these steps:

- On the Oracle server, open the command prompt and type the following command: `SQLPLUS /nolog`
- Type `connect<sys/password@SID>;`
- Type the following SQLPlus command:

```
SQLPLUS ALTER DATABASE RENAME FILE <old path from backup to  
any redolog file name> to <path to expected restored redolog  
file name>;
```

For example, ALTER DATABASE RENAME FILE
'D:\ORACLE\ORADATA\JACOB\REDO01.LOG' to
'C:\ORACLE\ORADATA\JACOB\REDO01.LOG';
- In the command prompt, type RMAN, and then type the following command in the RMAN prompt:
`alter database open resetlogs;`
- Close the command prompt.

The recovery is complete.

Requirements for recovering the complete Oracle instance and database using the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can use the same Oracle server for the recovery. You can also use these instructions when you configure a new physical server that uses the same server name and SID name.

To successfully complete the recovery using this scenario, you must have the following items:

Table H-9 Requirements when you recover using the original Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you log in.
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: ■ controlfile ■ datafiles ■ archive logs For Oracle 12c and later supported versions, you must perform a full backup of the container database.
The original Oracle server	To successfully recover the Oracle system in a disaster recovery scenario, you must restore to the original Oracle server.

Recovering the complete Oracle instance and database using the original Oracle server

You can use the same Oracle server for a recovery if you experience a complete loss, deletion, or destruction of the Oracle instance or database.

See [“Requirements for recovering the complete Oracle instance and database using the original Oracle server”](#) on page 1258.

To recover the complete Oracle instance or database using the original Oracle server

- 1 Recreate the Oracle database using the same name you used for the original database that was lost.
- 2 Find and rename the pwd<SID>.ora file.
- 3 Do the following in the order listed to create a new pwd<SID>.ora file:
 - Open a command prompt.
 - Type the following command:
`orapwd file=path\pwsid.ora password=<password>`
- 4 Type the following commands in the order listed:
 - `RMAN`
 - `CONNECT TARGET <sys/password@sid>;`
 - `SHUTDOWN ABORT;`
 - `STARTUP NOMOUNT;`
 - `SET DBID<dbid ID>;`
- 5 At the Backup Exec server, launch the Backup Exec Restore Wizard.
- 6 Select the appropriate ControlFile to restore.

The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.
- 7 After the restore job completes, exit Backup Exec.
- 8 At the Oracle server command prompt, type:
Alter database open resetlogs;
- 9 Close the command prompt.

Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can restore the instance and database to a computer other than the original Oracle server.

See [“Recovering the complete Oracle instance and database using the original Oracle server”](#) on page 1258.

To successfully complete the recovery, you must have the following items:

Table H-10 Requirements when you recover using a new or alternate Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you log in.
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: ■ controlfile ■ datafiles ■ archive logs

Recovering the complete Oracle instance or database to a computer other than the original Oracle server

You can restore an Oracle instance or database to a computer other than the original Oracle server.

See [“Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server”](#) on page 1259.

To recover the complete Oracle instance and database to a computer other than the original Oracle Server

- 1 Recreate the Oracle instance using the same name you used for the original instance that was lost.
- 2 Find and rename the pwd<SID>.ora file.
- 3 Do the following in the order listed to create a new pwd<SID>.ora file:
 - Open a command prompt.
 - Type the following command:

```
orapwd file=path\pwdsid.ora password=<password>
```
- 4 Type the following commands in the order listed:
 - RMAN
 - CONNECT TARGET <sys/password@sid>;
 - SHUTDOWN ABORT;
 - STARTUP NOMOUNT;

■ SET DBID<dbid ID>;

- 5 At the Backup Exec server, launch the Backup Exec Restore Wizard.
- 6 Select the appropriate ControlFile to restore.
- 7 Select the option to restore **To a different Oracle server**, and then select the appropriate options.
- 8 After the restore job completes, exit Backup Exec.

The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.
- 9 Move to the Oracle server.
- 10 Type **Alter database open resetlogs;**
- 11 Do one of the following:

If an error is encountered while Oracle tries to open the database	Note the online redo log path and then update the path.

If an error does not occur	Do nothing. The disaster recovery is complete.

Best practices for Backup Exec Agent for Oracle on Windows and Linux Servers

Best practices include tips and recommendations to help you use Backup Exec Agent for Oracle on Windows and Linux Servers (Oracle Agent) effectively. For more information about the Oracle Agent, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the Oracle Agent effectively:

- Enable the Oracle archive log mode and the Oracle automatic archival of log files.
- Know the DBID and other important configuration details of the database.
- Know the names of the init<SID>.ora and the spfiles for the instances on the Oracle server.

- Do not store the RMAN Repository on the same server that holds the database that you want to back up.
- Back up your current control file when you run a Database Administrator (DBA) initiated job. If you have a backup of the current control file, then you do not have to search media to find a control file that is available for recovery.
- Test recovery scenarios often to get comfortable with the restore procedures. Oracle recovery can be complex and is often time-sensitive due to the nature of the data involved. We recommend that you coordinate test plans and configuration activities with your Oracle DBA to be sure that restore procedures are confirmed.
- Use RMAN scripts to do the following:
 - Delete all archive log copies in a multiplexed archive log configuration.

Note: You can use Backup Exec to delete all non-multiplexed, single location archive logs.

- Run RMAN optimization.

The following best practices should be considered when you back up Oracle databases:

- Take a full backup whenever you make structural changes to a database.
- Do not delete archived log files unless you have two confirmed backups of each log.
- Create Oracle-specific media sets and backup jobs for the following reasons:
 - RMAN can manage media retention and can communicate to the Backup Exec server that backup sets are expired. RMAN can successfully manage the media's retention period as long as unrelated backup sets are not present. Unrelated backup sets may have retention periods that are longer than the RMAN retention period.
 - Media sets for Oracle backups should have a retention period that is greater than the `CONTROL_FILE_RECORD_KEEP_TIME` setting. By default, the `CONTROL_FILE_RECORD_KEEP_TIME` is 7 days. If the media sets for the Oracle backups have a greater retention period, backup sets are not overwritten and RMAN is not updated.
 - When you configure multiple job streams in Oracle, additional resources such as file systems can cause more devices than expected to allocate drives.

- Load balance Oracle jobs between managed Backup Exec servers in a CASO environment. However, this scenario means that archived log file backups may reside on multiple managed Backup Exec servers, which makes restores complicated.
- Consider port re-assignments when you use RALUS in a Linux environment. Applications such as Webmin that uses port 10000 can interfere with RALUS operations.
- Enable the Oracle block change tracking for faster incremental backups.
- Enable Backup Exec compression when you configure general options for backup jobs.
- Ensure that you enter the fully qualified domain name of the Oracle server when you add it to list of servers on the **Backup and Restore** tab.
- Ensure that you add the fully qualified domain name of the Oracle server and the logon account name to the Backup Exec server's list of Oracle servers and authentication credentials.

The following best practices must be considered if you use the Oracle 12c database:

- Take a full backup of a container database (CDB) whenever there are any structural changes, such as addition of a new pluggable database (PDB).
- Include the root of the CDB in the backups to ensure that metadata of the CDB is always backed up.
- Run the database in the archive log mode to ensure that the database can be recovered to point in time.
- If the CDB is in no-archive log mode, then before backing up the PDBs, shut down the CDB. To avoid shutting down the CDB, you can either run the database in archive log mode or run a DBA-initiated backup of PDBs.
- Oracle recommends users to not only restore the root because it might cause metadata inconsistencies. Instead, you should recover the whole CDB.
- If the PDB point-in-time (PIT) restore fails, then consecutive jobs might also fail with the following error message:
 ORA-19852: Error creating services for auxiliary instance.
 This error occurs because the previous failed PIT restore attempts were not cleaned up properly. To solve this issue, you must clean the failed database PIT restore attempts. Perform the following steps to clean up the failed auxiliary service creation attempts:
 - Use the Database (DB) PIT recovery package to clean up the auxiliary instance in case of failed PIT jobs:

```
SQL> exec dbms_backup_restore.manageauxinstance ('DBPITR',1);
```

- Also, run the last set of commands in the RMAN script to clear the RMAN configuration.

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' CLEAR;
```

```
CONFIGURE AUXILIARY CHANNEL DEVICE TYPE 'SBT_TAPE' CLEAR;
```

Backup Exec Agent for Enterprise Vault

This appendix includes the following topics:

- [About the Agent for Enterprise Vault](#)
- [Requirements for the Enterprise Vault Agent](#)
- [About installing the Enterprise Vault Agent](#)
- [About backup methods for Enterprise Vault backup jobs](#)
- [About backing up Enterprise Vault components](#)
- [About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)
- [Restoring Enterprise Vault](#)
- [Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database](#)
- [Best practices for the Enterprise Vault Agent](#)
- [About the Backup Exec Migrator for Enterprise Vault](#)
- [Enterprise Vault logon account](#)
- [Enterprise Vault options](#)

About the Agent for Enterprise Vault

The Backup Exec Agent for Enterprise Vault (Enterprise Vault Agent) is installed as part of the Agent for Applications and Databases.

The Enterprise Vault Agent provides data protection for the following Enterprise Vault components:

- Sites
- Vault Store Groups
- Databases
- Indexes
- Vault partitions

The Enterprise Vault Agent can help provide a disaster recovery solution for the data that is archived with Enterprise Vault. Recovery of the archived data is not dependent on the archive source, such as Exchange Server or a specific file system.

The Enterprise Vault Agent lets you do the following:

- Back up and restore Enterprise Vault archives from open or closed vault store partitions.
- Back up and restore individual Enterprise Vault vault store groups from an Enterprise Vault site.
- Back up and restore Enterprise Vault sites, databases, and index locations.

When you back up Enterprise Vault servers the following Enterprise Vault components can be backed up along with the vault partitions:

- Enterprise Vault Directory and Monitoring databases
- Enterprise Vault Audit, FSA Reporting, and Fingerprint databases
- Enterprise Vault vault store databases
- Enterprise Vault indexing files

If you install the Enterprise Vault Compliance and Discovery Accelerator products, the following components can be backed up:

- Enterprise Vault Compliance Accelerator and Discovery Accelerator Configuration databases
- Enterprise Vault Compliance and Discovery Accelerator Customer databases
- Enterprise Vault Discovery Accelerator Custodian database

The Enterprise Vault Agent uses Enterprise Vault Backup mode to back up Enterprise Vault components. By using the Backup mode, the Enterprise Vault Agent can back up Enterprise Vault components without having to suspend Enterprise Vault archiving operations.

For example, when you select a vault store group or site for backup, the individual vault store or indexes are placed in Backup mode. Backup mode lets Enterprise

Vault continue archiving operations in other vault store groups or sites. After the backup job successfully completes, the Enterprise Vault Agent takes the Enterprise Vault components out of Backup mode so that those components can continue archival operations.

While Enterprise Vault versions 8.x, 9.x, and 10.x all implement Backup mode, Enterprise Vault 9.x and 10.x offer you more flexibility with your vault store backup jobs. With Enterprise Vault 9.x and 10.x, you can run multiple backup jobs of the same Enterprise Vault 9.x and 10.x vault store simultaneously. With Enterprise Vault 8.x, multiple vault store backup jobs must run one at a time.

For example, you can create multiple backup jobs to back up a vault store. Each backup job includes in its selection list one or more unique partitions of the vault store. Under Enterprise Vault 9.x and 10.x, the partitions are backed up simultaneously when the different backup jobs access them at the same time. Under Enterprise Vault 8.x, the partitions are backed up in both backup jobs; however they are backed up sequentially. The first backup job must finish before the second job starts, or else a backup job failure occurs.

Note: With all versions of Enterprise Vault, the Enterprise Vault Agent automatically backs up the vault store database whenever an open partition is backed up.

The Enterprise Vault Agent backs up the Compliance Accelerator and Discovery Accelerator application databases while they are online. It does not place the databases in Read-only mode or Backup mode before it backs them up.

The Enterprise Vault Agent runs a physical check on each Enterprise Vault database before it backs them up. The Enterprise Vault Agent also runs a physical check on each database before you restore them.

Note: The Enterprise Vault Agent uses physical database consistency checks because physical checks consume less system resources than other types of Database Consistency Checking options.

Backing up and restoring Enterprise Vault databases and related components require specific user account credentials for each Enterprise Vault component you protect.

Table I-1 Supported user accounts that are required to back up and restore Enterprise Vault components

Enterprise Vault components	User credentials
Enterprise Vault databases and components (vault store, indexes, partitions, vault store database, Directory, Monitoring, Fingerprint, FSA Reporting, and Audit databases)	The following credentials are required: ■ Vault Service account ■ Domain Admin account with Role Based Admin privileges You can also use any domain user account that meets the following requirements: ■ The user account must be included in the Administrator group on all servers where Enterprise Vault partitions and Enterprise Vault databases reside. ■ The user account must have Backup-related Role Based Admin privileges for the vault store and the index location. Backup-related Role Based Admin privileges include: ■ EVT Manage Vault Store Backup Mode ■ EVT Manage Index Location Backup Mode To configure Role Based Admin privileges for a Windows Domain Admin account, see your Enterprise Vault documentation.
Compliance Accelerator and Discovery Accelerator	The following credentials are required: ■ Vault Service account ■ Domain Admin account A user account that is a member of the Administrator's group on the computers where the Compliance and Discovery Accelerator applications reside.

When you back up specific Enterprise Vault components, the other Enterprise Vault components are automatically backed up at the same time. Backup Exec includes these components to hasten an Enterprise Vault recovery.

Table I-2 Enterprise Vault databases that are automatically backed up

When you back up this	Backup Exec automatically backs up this	Description
Enterprise Vault site	Directory database	Backup Exec automatically backs up the Directory database that is associated with the Enterprise Vault site.
Open partition	Vault store database	Backup Exec automatically backs up the vault store database that is associated with the open partition.

Over time the amount of data that Enterprise Vault stores continues to grow. At some point, you may observe that as the data moves through its usage lifecycle, you no longer access it as frequently. You can use the Backup Exec Migrator for Enterprise Vault to automatically migrate the older Enterprise Vault data to the storage devices that Backup Exec manages.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1287.

Requirements for the Enterprise Vault Agent

Review the following requirements before you use the Agent for Enterprise Vault (Enterprise Vault Agent).

- You must create at least one partition on an Enterprise Vault server before the Enterprise Vault server can publish itself to Backup Exec.
- You must install the Backup Exec Agent for Windows and license the Enterprise Vault Agent on any computer that hosts an Enterprise Vault component.

Note: The Enterprise Vault Agent uses the Agent for Windows to back up all NTFS shares on a remote computer that contains Enterprise Vault data. However, if the Agent for Windows is not installed, the Enterprise Vault Agent uses Microsoft's Common Internet File System (CIFS) to back up the data.

For a device or a filer that does not support the Agent for Windows, the Enterprise Vault Agent uses CIFS to back up the data. It is recommended that you create separate backup jobs when you want to do NDMP backups of Enterprise Vault data. You may see a significant performance improvement of NDMP backups with the Backup Exec NDMP feature.

About installing the Enterprise Vault Agent

The Agent for Enterprise Vault (Enterprise Vault Agent) is installed as part of the Agent for Applications and Databases feature. To back up all Enterprise Vault servers, the Enterprise Vault Agent must be installed on each Enterprise Vault server in your environment. In addition, the Enterprise Vault Agent must also be installed on any remote computer where Enterprise Vault components are installed. If the Compliance and Discovery Accelerators are installed on remote computers, the Enterprise Vault Agent must be installed on those computers too.

You can install the Enterprise Vault Agent in the following ways:

- Install it automatically from the Backup Exec server as part of an Agent for Windows installation to the local Enterprise Vault server. After you finish the installation, you may need to configure the Enterprise Vault Agent to publish itself to a Backup Exec server of your choice.
See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 958.
- Install the required Enterprise Vault Agent licenses on the Backup Exec server. After you install the licenses, you can push-install the Backup Exec Agent for Windows to all Enterprise Vault servers and the computers where other Enterprise Vault components are installed.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 68.

About backup methods for Enterprise Vault backup jobs

You can select a backup method that depends on the Enterprise Vault object that you want to back up.

The following table describes the type of Enterprise Vault backup jobs you can run. The table also describes the backup methods that are available for each type of backup job.

Table I-3 Backup methods to use with Enterprise Vault backup jobs

To back up	Select	Description
Directory and Monitoring databases Audit database and FSA Reporting database	Full, differential, or incremental backup method	Directory, Monitoring, Audit, and FSA Reporting database backups can use the full and incremental backup methods. These databases cannot be backed up using the differential backup method. If you select the differential backup method, Backup Exec does a full backup instead . Note: Selecting an incremental backup method backs up the database transaction logs and truncates them.
Vault database and Fingerprint database	Full, differential, or incremental backup method	Vault database and Fingerprint database backups can use all three backup methods: Full, differential, and incremental. Note: Selecting an incremental backup method backs up the database transaction logs and truncates them.
Vault partitions and index locations	Full, differential, or incremental backup methods	You can use all of the backup methods that are available for standard file system backup jobs.

When you combine Enterprise Vault components in a backup job, each component may use a backup method that differs from what you selected for the overall job. For example, you create a job that uses the differential backup method to back up both a Directory database and a partition. However, because a Directory database cannot be backed up using the differential method, Backup Exec uses the full backup method to back up the Directory database. This results in fast

and easy restores. After the Directory database is backed up, Backup Exec uses the differential backup method to back up the partition.

Use the following table as a guide.

Table I-4 Actual backup methods that are used for Enterprise Vault components

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Directory and Monitoring databases	F	F	I Always truncates the transaction logs
Vault store database	F	D	I Always truncates the transaction logs
Audit database	F	F	I Always truncates the transaction logs
FSAReporting database	F	F	I Always truncates the transaction logs
Fingerprint database	F	D	I Always truncates the transaction logs
Partition	F	D	I
Index root path	F	D	I

Table I-4 Actual backup methods that are used for Enterprise Vault components (*continued*)

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Compliance Accelerator/Discovery Accelerator Configuration database Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with run-time versions of Enterprise Vault.	F	F	I Always truncates the transaction logs
Compliance Accelerator/Discovery Accelerator Customer database Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with run-time versions of Enterprise Vault.	F	D	I Always truncates the transaction logs

Table I-4 Actual backup methods that are used for Enterprise Vault components (*continued*)

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Discovery Accelerator Custodian database Note: Also includes the Discovery Accelerator Custodian databases that are installed with run-time versions of Enterprise Vault.	F	D	I Always truncates the transaction logs

See [“Backup methods in Backup Exec”](#) on page 189.

See [“About backing up Enterprise Vault components”](#) on page 1274.

Enterprise Vault backup options

You can select a backup method that is based on the type of Enterprise Vault database that you want to back up.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 1271.

See [“Backing up data”](#) on page 158.

See [“About backing up Enterprise Vault components”](#) on page 1274.

About backing up Enterprise Vault components

You can select any or all of the Enterprise Vault components for backup when you create a backup job. If you select all of the components for backup in the same job, recovery time is faster. However, if you create multiple backup jobs for the components, the backup jobs run faster.

The Enterprise Vault components that you can select are described in the following table, along with recommendations for backup:

Table I-5 Enterprise Vault components

Enterprise Vault Component	Description
Directory database	The Directory database is a Microsoft SQL Server database that contains configuration data. After the database is populated, the amount of data in the Directory database changes very little over time. You should back up the Directory database after you add or remove any Enterprise Vault component. You should also back up the Directory database if you change the location of any component. Configuration changes can include creating vault stores, creating vault store partitions, and changing vault store partition statuses.
Monitoring database	Enterprise Vault includes a Monitoring agent on each Enterprise Vault server. The Monitoring agent monitors the following: ■ The status of Enterprise Vault services and tasks. ■ Performance counters for vault stores, disk space, memory, and processors. ■ The status of Exchange Server journal mailbox target archiving targets, including item counts for Inbox, Archive Pending, and failed operations such as Failed DL Expansion. The Monitoring agent collects monitoring data at scheduled intervals, typically every few minutes. All of the information that the Monitoring agent collects is stored in a Microsoft SQL Server database called the Monitoring database.
Fingerprint databases	The Fingerprint databases contain the single instance storage-related information for all of the vault stores in the vault store group. If you enable single instance storage of archived items, you should back up the Fingerprint databases on a regular basis.

Table I-5 Enterprise Vault components (*continued*)

Enterprise Vault Component	Description
Index location	The index location stores all of the archived data content that is indexed to enable fast searching and retrieval of archived items. The indexing data is stored in index files in the location that is specified when you install Enterprise Vault. You should back up the index location on a regular basis.
Vault store group	The vault store group is a logical entity. If you select it for backup, all of the vault databases, vault store partitions, and the Fingerprint databases are backed up. Because these components are closely related, you should consider selecting the vault store group to back up all of these components together.
Vault store	The vault store is a logical entity. If you select it for backup, all of the vault databases and the vault store partitions are backed up.
All partitions	A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. If you select All Partitions for backup, then all of the vault store partitions in the vault store are selected for backup. Note: When you back up an open partition, Backup Exec automatically backs up the vault store database.
Site	An Enterprise Vault site is a logical representation of an installation of the Enterprise Vault. If you select this component for backup, the Directory database is also automatically backed up.
Compliance Accelerator database and Discovery Accelerator database	These databases are installed as optional add-ons to Enterprise Vault and are part of the Discovery Accelerator and Compliance Accelerator products.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 205.

See [“Adding a stage to a backup definition”](#) on page 220.

About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases

Backup Exec automatically checks the physical consistency of an Enterprise Vault database before a backup job and after a restore job. It also checks the consistency of the Compliance and Discovery databases before a backup job and after a restore job. Backup Exec uses Microsoft SQL Server's Physical Check Only utility for consistency checks of the databases. In the event a consistency check fails, Backup Exec continues with the job and reports the consistency check failures in the Backup Exec job log.

For more information about the Physical Check Only utility, see the Microsoft SQL Server documentation.

Restoring Enterprise Vault

Review the following before you begin an Enterprise Vault restore operation.

- When you restore an Enterprise Vault installation, you should restore the Directory database in a separate restore job. After you successfully restore the Directory database, you can restore other Enterprise Vault components and partitions.
- When you restore Enterprise Vault databases, you can select the options that either leave databases in a ready-to-use state or in a non-operational state. The non-operational state options that you select apply to all Enterprise Vault databases except the vault store database. When you restore an Enterprise Vault vault store database, the Agent for Enterprise Vault (Enterprise Vault Agent) places the vault store database in Enterprise Vault Backup mode. If the vault store database remains in a non-operational state after the restore job completes, the Enterprise Vault Agent cannot remove it from Backup mode.
- If you select the option that leaves the databases ready to use, the Enterprise Vault Agent restores the vault store database in a ready-to-use, operational state. The vault store database's operational status is maintained even when you select additional backup sets for restore in the same vault

store database restore job. Additional backup sets can include Full, Differential, and Incremental backup methods.

- If you select the option that leaves the databases in a nonoperational state, the Enterprise Vault Agent prompts you to stop the **Enterprise Vault Storage Service** before you start the vault store database restore operation. You can restart the vault store restore operation again after the Enterprise Vault Storage Service stops.

As a best practice, it is recommended that you restore the vault store database in a ready-to-use state. When you restore the vault store database in a nonoperational state, Enterprise Vault cannot remove it from Backup mode after the restore operation finishes.

See [“Enterprise Vault restore options”](#) on page 1279.

- You can individually restore Enterprise Vault components. Before you begin the restore, the databases and other components may or may not exist on the destination Enterprise Vault server. If the databases do not exist, you can restore them using the Enterprise Vault Agent. After the restore job completes, you must configure Enterprise Vault to use the restored databases. To configure Enterprise Vault to use the restored databases, see your Enterprise Vault documentation.

These items include the following:

- Enterprise Vault 8.x/9.x/10.x Directory, Monitoring, Audit, FSAReporting, and Fingerprint databases
- Vault store databases, indexes, and partitions.
- Compliance and Discovery Accelerator Configuration and Customer databases.
- Discovery Accelerator Custodian database
- It is recommended that you use the Enterprise Vault service account or an account with rights to access the restore selections as the default logon account. Otherwise, you may have to enter proper credentials for each Enterprise Vault resource that you select for restore.
- After you restore Enterprise Vault, a message appears that says you need to run Enterprise Vault recovery tools. The recovery tools are used to re-synchronize Enterprise Vault with the newly restored databases after you complete the restore.

For information on running the Enterprise Vault recovery tools, see your Enterprise Vault documentation.

Before you restore Enterprise Vault sites, servers or other components, you should have the following items installed on the destination computer:

- Enterprise Vault
- The Backup Exec Agent for Windows

Note: You must install the Agent for Windows on remote Enterprise Vault computers where you want to restore Enterprise Vault components.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

Enterprise Vault restore options

Use the following table to select the restore option you want to use when you restore the Enterprise Vault databases.

Table I-6 Enterprise Vault restore options

Item	Description
Automatically take the Enterprise Vault databases offline when restoring selected databases	Takes the shared Enterprise Vault Directory, Monitoring, Audit, FSA Reporting, and Fingerprint databases offline so Backup Exec can replace them during a restore job. Note: If you don't use this option, you must stop the Directory and Admin services on all Enterprise Vault servers before you restore the previously mentioned databases. In addition, you must also stop the Accelerator Manager server on all of the Compliance Accelerator servers and the Discovery servers. Only after you stop the Accelerator Manager can you restore the Customer, Configuration, and Custodian databases. It also terminates connections to the following: ■ Monitoring database ■ Audit, Fingerprint, and FSA Reporting databases (Enterprise Vault 8.x, 9.x, 10.x only) ■ Configuration, Customer, and Custodian databases When the restore job completes, you must manually restart the Enterprise Vault Admin and Directory services on your Enterprise Vault server. After you restart the services, the services reconnect to the restored databases and Enterprise Vault begins archival operations again. Note: This option causes the Enterprise Vault Admin and Directory services on all Enterprise Vault servers to terminate their connections to the Directory database that you restore. It also terminates the connections to the Enterprise Vault Accelerator Manager database.

Table I-6 Enterprise Vault restore options (*continued*)

Item	Description
Do not take the Enterprise Vault databases offline	Leaves all Enterprise Vault databases online. If you use this option, you must stop the Directory and Admin services on all Enterprise Vault servers before you restore the previously mentioned databases. In addition, you must also stop the Accelerator Manager server on all of the Compliance Accelerator servers and the Discovery servers. Only after you stop the Accelerator Manager can you restore the Customer, Configuration, and Custodian databases.
Leave the database ready to use; additional transaction logs or differential backups cannot be restored	Rolls back all uncompleted transactions when you restore the last database, differential, or log backup. After the recovery operation, the database is ready for use. If you do not select this option, the database is left in an intermediate state and is not usable. If you select this option, you cannot continue to restore backups. You must restart the restore operation from the beginning.
Leave the database nonoperational; additional transaction logs or differential backups can be restored	Creates and maintains a standby database. By using this option, you can continue restoring other backup sets for non-operational databases. See your SQL documentation for information on standby databases.

Note: It is recommended that you select all required backup sets when you run a single restore job for a vault store database. All required backup sets can include full, differential, and incremental backup sets. The vault store database should also be restored in a ready-to-use state after the restore job completes.


See [“Methods for restoring data in Backup Exec”](#) on page 235.

About restoring individual files and folders with the Enterprise Vault Agent

The Agent for Enterprise Vault (Enterprise Vault Agent) supports individual file and folder restores from vault store partition backups. You can also restore complete index locations or individual folders from Enterprise Vault index backups.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

Figure I-1 Restoring individual files from vault store partitions and complete folders from an Enterprise Vault index



About automatic redirection of Enterprise Vault components under an Enterprise Vault server

You can change the location of the vault store databases, Fingerprint databases, or partitions to a location that differs from where they were backed up. During restores of the vault store database, Fingerprint databases, or partitions, the Agent

for Enterprise Vault (Enterprise Vault Agent) detects the location change. It then automatically redirects the component restores to the new location.

Note: Automatic redirected restores of the vault databases, partitions, or Fingerprint databases occur when you change only the location of these Enterprise Vault components. The names of the partitions, vault stores, and vault store groups must not change from the time the partition was originally backed up.

See [“Restoring Enterprise Vault”](#) on page 1277.

Redirecting a restore for an Enterprise Vault component

You can redirect the restore of the Enterprise Vault components.

The following table describes the requirements for redirecting a restore for an Enterprise Vault component:

Table I-7 Requirements for redirecting a restore for an Enterprise Vault component

Component	Requirement
All Enterprise Vault components	The following are requirements for redirecting the restore of all Enterprise Vault components: ■ All Enterprise Vault components must already exist on the server to which you redirect the restore. If they do not exist, you must create them. See your Enterprise Vault documentation. ■ The Backup Exec logon account that you use must have the same credentials as the Vault Store service account.
Enterprise Vault databases	The following are requirements for redirecting the restore of the databases: ■ You must create a separate job for each database that you want to redirect. ■ You must redirect all databases to the same SQL server.

Table I-7 Requirements for redirecting a restore for an Enterprise Vault component *(continued)*

Component	Requirement
Vault store databases	The following are additional requirements for redirecting the restore of a vault store database: ■ Ensure that the Directory database already exists on the server to which you redirect the restore. ■ Ensure that the Directory database contains an entry for the vault store that uses the new SQL Server name.
Vault store partition	The following are additional requirements for redirecting the restore of a vault store partition: ■ The vault store must already contain a vault partition with the same name. If a vault partition does not exist, you must create it.
Index location	The Directory database must already be configured with the new index location. See your Enterprise Vault documentation.

Redirection options for Enterprise Vault

You can redirect a restore job for Enterprise Vault components.

See [“Redirecting a restore for an Enterprise Vault component”](#) on page 1283.

Table I-8 Redirection options for Enterprise Vault

Item	Description
To a new Microsoft SQL server	Redirects the restore jobs of Enterprise Vault databases and Accelerator databases to a different SQL Server. Displays the name of the server to which you want to redirect the restore job for a vault store. Note: Vault store databases are restored for Enterprise Vault 8.x, 9.x and 10.x only.

Table I-8 Redirection options for Enterprise Vault *(continued)*

Item	Description
Instance	Displays the name of the instance of the SQL Server to which you want to redirect the restore job for a vault store.
Restore index root to a new location	Redirects the restore job for the index root to a new location. If you redirect the restore of the Enterprise Vault server, you can specify an alternate path on the destination server. You can also redirect the index root location to an alternate path on the original server.
Path	Displays the path name to which you want to redirect the restore job for an index root.
Restore partition root to a new location	Redirects the restore job for a vault store partition to a new location. Partitions are restored for Enterprise Vault 8.x, 9.x , and 10.x only.
Path	Displays the path name to which you want to redirect the restore job for a vault store partition.
Enterprise Vault logon account	Specifies the logon account to use.

Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database

Use the following steps to configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database.

To configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database

- 1 At each Enterprise Vault server, use Enterprise Vault to change the name of the previous SQL Server computer. Change the name to the name of the SQL Server computer that now holds the Directory database.

See your Enterprise Vault documentation.

- 2 Restart the Enterprise Vault Admin service on all Enterprise Vault servers that use the Directory database.

Two directory names appear in the backup selections view after you restart the Enterprise Vault Admin service on the Enterprise Vault server.

For example, **Directory on <OldSQL_computer_name>** and **Directory on <NewSQL_computer_name>**.

- 3 On the **Backup and Restore** tab, right-click the Enterprise Vault server that you want to back up.
- 4 On the **Backup** menu, select the backup definition that you want to use.
- 5 In the **Selections** box, click **Edit**.
- 6 Expand **Directory on <SQL server computer where you moved the Directory database>**.
- 7 Expand all items under **Directory on <SQL server computer where you moved the Directory database>**.

The **Directory** and **Monitoring** databases, Enterprise Vault 8.x, 9.x, or 10.x **FSA Reporting** and **Audit** databases, and the Enterprise Vault sites should appear. In addition, the Directory database should display the new SQL Server name and instance where it was redirected.

When you configure a new Directory database backup job, you must select the Directory database from the current Directory server. Backup Exec automatically removes the previous Directory server name 13 days after you complete the Directory database move.

- 8 To manually remove the previous server name, right-click **Directory on <OldSQL_computer_name>**.
- 9 Click **Remove**.

See your Enterprise Vault documentation.

Best practices for the Enterprise Vault Agent

It is recommended the following best practices when you use the Agent for Enterprise Vault (Enterprise Vault Agent).

- Back up the Enterprise Vault Directory database after you make any configuration changes in Enterprise Vault.
- Restore the Enterprise Vault Directory database in a separate Backup Exec restore job.
- Restore all Full, Differential, and Incremental backup sets of the vault store database in a single restore job.
- Do not allow the backup window and archive window to overlap.
- Do not allow the backup window and the migration window to overlap.
- Make sure Enterprise Vault components are not in Backup mode before you back up the Enterprise Vault Directory database.
- If you install both the Backup Exec NDMP feature and the Enterprise Vault Agent, pick only one product to protect an Enterprise Vault partition that resides on NDMP filers.
- Do not change the recovery model of any database that is created by Enterprise Vault. Enterprise Vault configures each database in full recovery mode when it creates them.

For more information about the best practices to use Backup Exec Agent for Enterprise Vault and the Backup Exec Migrator, refer to *Backup Exec Best Practices*.

See [“About the Agent for Enterprise Vault”](#) on page 1265.

About the Backup Exec Migrator for Enterprise Vault

The Backup Exec Migrator for Enterprise Vault (Backup Exec Migrator) lets you automatically migrate archived Enterprise Vault data to the storage devices that Backup Exec manages. By migrating the archived Enterprise Vault data from a partition, you can reclaim disk space on the Enterprise Vault server without incurring the cost of additional hardware.

By migrating Enterprise Vault archive data to the Backup Exec server storage devices, you also ensure an added level of storage redundancy using an off-host environment.

See [“How the Backup Exec Migrator works”](#) on page 1288.

See [“Configuring the Backup Exec Migrator”](#) on page 1297.

Backup Exec Migrator for Enterprise Vault requirements

- Before you configure the Backup Exec Migrator, ensure that your Enterprise Vault server meets the following requirements:
- Backup Exec Agent for Enterprise Vault must be installed on the Enterprise Vault server.
 - Enterprise Vault migration and collections must be enabled for the Enterprise Vault partition from which you want to migrate data.
 - Enterprise Vault 8.0 SP3 or higher must be installed on the Enterprise Vault server.

How the Backup Exec Migrator works

Enterprise Vault automatically initiates all data migration operations from the Enterprise Vault server after you configure the Backup Exec Migrator. Enterprise Vault makes decisions on what should be migrated based on the archival policies and the data retention policies that you configure in the Enterprise Vault Administration Console. The Backup Exec Migrator then migrates the archived data to a Backup Exec server after Enterprise Vault collects the eligible data from the vault store partitions. When you configure migration options for a partition, you can set the migration period. All migration options are configured at the Enterprise Vault server.

Table I-9 Enterprise Vault data migration process

Action	Notes
Enterprise Vault archives eligible partition data that is based on the file size or the file creation date.	All data that is eligible for archive is determined in the partition where you want to migrate data. See your Enterprise Vault documentation.

Table I-9 Enterprise Vault data migration process (continued)

Action	Notes
After Enterprise Vault completes the archival process, an Enterprise Vault collection process collects the archived data.	The collection process places the archived data into Windows .cab files. The .cab files are stored in the partition where the migration occurs. Eligible data can include Enterprise Vault files with the following extensions: ■ .dvf ■ .dvssp ■ .dvsc ■ .dvs Note: Some eligible data cannot be compressed into .cab files due to file size restrictions. However, the Backup Exec Migrator still migrates the data during the migration operation. See your Enterprise Vault documentation.

Table I-9 Enterprise Vault data migration process (continued)

Action	Notes
The Backup Exec Migrator initiates the migration of the archived data files to a Backup Exec server.	

Table I-9 Enterprise Vault data migration process (*continued*)

Action	Notes
	Migration period schedules are determined when you configure migration for a partition and when you configure a collection schedule for the partition. See “Configuring Enterprise Vault collections” on page 1298. See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1302. If you follow the configuration recommendations for the Backup Exec Migrator and Enterprise Vault partitions, one migration job for each partition runs during a migration period. However, the Backup Exec Migrator may create separate migration jobs for each partition folder if you do not follow the configuration recommendations. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance. Note: If you schedule a file retrieval request from the Enterprise Vault server between migration periods, separate jobs are created even though you followed the configuration recommendations. In this case, the Backup Exec Migrator automatically creates separate jobs to facilitate retrieval of the requested file. During a migration operation, the restore job can be scheduled to run between migration jobs. If you do not follow the configuration recommendations, file retrieval performance can be affected. To ensure the most efficient migration and retrieval performance possible, follow the recommendations when you configure the Backup Exec Migrator and the Enterprise Vault partitions.

Table I-9 Enterprise Vault data migration process (*continued*)

Action	Notes
	See “Configuring the Backup Exec Migrator” on page 1297.
Backup Exec completes the migration process by moving all of the migrated files to storage devices.	It is recommended configuring two storage devices for staged migration operations. See “About using staged migrations with Backup Exec and the Backup Exec Migrator” on page 1292. See “Configuring the Backup Exec Migrator” on page 1297.

After Backup Exec migrates the .cab files to the storage devices, you can review migration details by looking at the job history details for each Enterprise Vault server where the migration occurs.

See [“About the Job History”](#) on page 274.

About using staged migrations with Backup Exec and the Backup Exec Migrator

When you configure Backup Exec to work with the Backup Exec Migrator, it is recommended that you configure two storage devices for staged migration operations. When you consider the devices to use, consider selecting a high performance disk-based storage and a slower performance tape device. By using two devices, archived data can be migrated in two stages.

During the first stage, Backup Exec migrates the data it receives from the Backup Exec Migrator to a disk-based storage on a high performance hard drive. By using a disk-based storage, you can minimize the amount of time it takes to perform the initial migration. During the second migration stage, Backup Exec creates a duplicate job to migrate the archived data from the disk-based storage to a tape device. You can schedule the duplicate job to move the archived data to a tape device at times when Backup Exec server activity is low.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1299.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1302.

About Backup Exec Migrator events

The Backup Exec Migrator generates events that specify the status of the tasks that it runs. The events also provide useful information for troubleshooting purposes. You can view the events on the computer where you installed the Enterprise Vault Storage Service by viewing the Windows Event Viewer. From the Event Viewer, you can see the events under **Enterprise Vault**. You can also view the events in the Enterprise Vault Dtrace Utility.

For more information on the Enterprise Vault Dtrace Utility, see your Enterprise Vault documentation.

See [“About Backup Exec Migrator logs”](#) on page 1293.

About Backup Exec Migrator logs

The Backup Exec Migrator can create log files that log all migration activity. The log files reside on both the Enterprise Vault server and the Backup Exec server. Backup Exec Migrator log files can help you troubleshoot migration issues.

Before you can view the log files, you must enable Backup Exec Migrator logging on the Enterprise Vault server and on the Backup Exec server. To enable Backup Exec Migrator logs on the Enterprise Vault server, edit the Windows registry.

See [“How to enable Backup Exec Migrator logging”](#) on page 1294.

You must also enable Backup Exec Migrator logging on the Backup Exec server.

See [“Using the Backup Exec Debug Monitor for troubleshooting”](#) on page 875.

Note: Partition Recovery Utility log files are enabled by default.

After you enable logging on the Enterprise Vault server and on the Backup Exec server, the following types of log files are created:

- VxBSA log files
For example, <computer_name>-vxbsa<00>.log
- Partition Recovery Utility log files
For example, partitionrecovery<00>.log
- Backup Exec server log files
For example, <computer_name>-bengine<00>.log

Each time the Backup Exec Migrator is started, separate VxBSA log files are created. As a result, each new log file's sequential number increments by one.

For example, <computer_name>vxbsa00.log, <computer_name>vxbsa01.log.

Similarly, a new log file is created each time the Partition Recovery Utility is started. As a result, each new Partition Recovery Utility log file's sequential number increments by one.

For example, `partitionrecovery00.log`, `partitionrecovery01.log`

Backup Exec server log file numbers also increment by one as multiple log files are created.

For example, `<computer_name>-bengine00.log`, `<computer_name>-bengine01.log`

You can find the log files in the following locations.

Table I-10 Backup Exec Migrator and Partition Recovery Utility log file locations

Log file	Computer	Directory location
VxBSA log files Partition Recovery Utility log files	Enterprise Vault server	C:<Backup Exec install path>\BACKUP EXEC\RAWS\logs
Backup Exec server log files	Backup Exec server	C:<Backup Exec install path>\Backup Exec\Logs

See [“About Backup Exec Migrator events”](#) on page 1293.

How to enable Backup Exec Migrator logging

Perform the following steps to enable Backup Exec Migrator VxBSA logging for the Backup Exec remote agent on an Enterprise Vault server.

Warning: Incorrect use of the Windows registry editor may prevent the operating system from functioning properly. Take great care when making changes to the Windows registry. Registry modifications should only be carried out by persons who are experienced in using the registry editor application. To ensure data continuity, make a complete backup of the registry and your system prior to making any registry changes.

1. On the Enterprise Vault server, click **Start**, then select **Run**.
2. Type `regedit` and press `Enter`.
3. In the Registry Editor, navigate to **My Computer**, then to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > Backup Exec for Windows > Backup Exec > Debug**.

4. Double-click the **VXBSAlevel** key.
5. Change the **Value** data field to 6. This enables full verbose logging.
6. Click **OK**.
7. Enable logging for the remote agent by setting the following registry value to 1: **HKLM\SOFTWARE\Veritas\Backup Exec For Windows\Backup Exec\Engine\Logging\CreateDebugLog**.
8. Close the Windows Registry Editor.

On the Backup Exec media server, enable logging for the engine service by setting the following registry keys:

- Set the **HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\Backup Exec For Windows\Backup Exec\Debug\Enabled** key to 1.
- Set the **HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\Backup Exec For Windows\Backup Exec\Debug\Verbose** key to 6.

After enabling Backup Exec Migrator logging, the log files can be found in the following directory on the Enterprise Vault server: **C:<Backup Exec install path>\BACKUP EXEC\RAWS\logs**

The log files generated are shown below:

- **ComputerName-vxbsa<XX>.log**
- **ComputerName-bengine<XX>.log** (If the Backup Exec media server is installed on the Enterprise Vault server.)
- **ComputerName-beremote<XX>.log**
- **ComputerName-TAOSrv<XX>.log**
- **ComputerName-TAOCli<XX>.log**

To enable the logging of events related to the Migrator on the Enterprise Vault server application, perform the following steps on the Enterprise Vault server:

1. Open a command prompt and navigate to the directory under which the Enterprise Vault is installed. This directory contains the **Dtrace.exe** file.
2. Run the **Dtrace.exe** application.
3. Enable verbose logging on the processes **StorageFileWatch** and **EVStgOfflineOpns** by executing the following commands:

- `set StorageFileWatch verbose`
- `set EVStgOfflineOpns verbose`

4. Enter the command `View` to view a list of processes for which you can enable verbose logging.
5. Set the log file by using the command `Log [log_file_name]`
6. To enable monitoring inside the console itself, use the command `mon`
7. Execute the Migrator tasks for which logs are required.
8. Press `ctrl+c` to exit the monitoring phase.
9. Disable logging to ensure that all log entries are stored in the log file. Use the command `log` to disable logging. You can then collect the log file with the "`log_file_name`" that you specified in the `Log` command earlier.

About deleting files migrated by Backup Exec Migrator

Enterprise Vault automatically deletes archived items when the item's Enterprise Vault retention periods expire. An Enterprise Vault retention period indicates how long Enterprise Vault retains archived items before it deletes them.

The Backup Exec Migrator maintains existing Enterprise Vault retention periods for archived items when it migrates the archived items to tape. As a result, when an item's data retention period expires, Enterprise Vault issues a command to delete the item from the storage tape that Backup Exec manages. To delete the expired archive item, the `.cab` file it resides in must be deleted from tape.

Note: Although the Backup Exec Migrator maintains existing Enterprise Vault retention periods, it does not initiate the deletion of expired archived items or archived partitions from tape. Only Enterprise Vault can initiate the deletion of expired items and partitions.

For more information on deleting expired items, see your Enterprise Vault documentation.

Because the `.cab` files may contain archived items with different retention periods, an expired item may be marked as deleted in the Backup Exec catalogs. However, it may not be immediately deleted from tape. All archived items in a `.cab` file must have expired retention periods before Enterprise Vault issues a command to delete the `.cab` file from tape.

Enterprise Vault can also delete entire archived vault store partitions from tape. After you delete an active Enterprise Vault vault store partition by using the Enterprise Vault Administration Console, Enterprise Vault deletes the associated archived partition from tape.

Backup Exec automatically recycles the tapes when all of the items on the tape are marked as deleted in the catalogs. Backup Exec checks for expired Enterprise Vault Migrator media once every 24 hours. If Backup Exec detects such media, it logically moves the media to the Scratch Media set and then generates an information alert informing you of the move.

Note: Expired Enterprise Vault Migrator media is defined as media that contains only migrated Enterprise Vault data that is marked as deleted in the Backup Exec catalogs.

See [“Managing tapes”](#) on page 491.

Note: You should ensure that migrated Enterprise Vault data remains accessible on the tapes that are used for migration purposes until the Enterprise Vault data retention periods expire. Therefore, it is recommended that you configure a retention period of 999 years for all tapes that are used for migration purposes.

See [“Overwrite protection periods and append periods in media sets”](#) on page 495.

Configuring the Backup Exec Migrator

All of the program files that are required to run the Backup Exec Migrator are installed when you install the Agent for Enterprise Vault (Enterprise Vault Agent) on the Enterprise Vault server. However, before you can use the Backup Exec Migrator, you must configure it to work with both a destination Backup Exec server and the Enterprise Vault server.

Table I-II Enterprise Vault configuration process

Step	Description
Step 1	Configure Enterprise Vault collections. See “Configuring Enterprise Vault collections” on page 1298.
Step 2	Configure the Backup Exec Migrator to work with a Backup Exec server. See “Configuring the Backup Exec Migrator to work with a Backup Exec server” on page 1299.

Table I-11 Enterprise Vault configuration process (*continued*)

Step	Description
Step 3	Configure the Backup Exec Migrator to work with Enterprise Vault. See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1302.

Use the following configuration recommendations for both the Backup Exec Migrator and the Enterprise Vault partitions:

- Configure the Enterprise Vault partitions to save migrated data locally. Do not configure Enterprise Vault partitions to delete files immediately after a migration operation finishes. See your Enterprise Vault documentation for details on configuring a partition for migration.
- Configure the Backup Exec server template to run staged migrations. See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1292.

Failure to follow the configuration recommendations results in degraded migration and retrieval performance.

Configuring Enterprise Vault collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault first needs to collect the data.

To configure Enterprise Vault collections

- 1 From the Enterprise Vault Console, navigate to a vault store partition from which you want to migrate data.
- 2 Right-click the partiition, and then click **Properties**.
- 3 On the **Collections** tab, check **Use collection files**.
- 4 Set collection options as appropriate.
See [“Vault store partition properties - Collections”](#) on page 1298.
- 5 Click **OK**.

Vault store partition properties - Collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault needs to collect the data to be migrated.

See [“Configuring Enterprise Vault collections”](#) on page 1298.

Table I-12 Vault store partition properties - Collection options

Item	Description
Use collection details	Lets you set Enterprise Vault as the collector.
Start at	Indicates the local time at which you want collection to start.
End at	Indicates the local time at which you want collection to finish. Enterprise Vault stops collecting at this time or when it has no more files to collect, whichever comes first.
Limit collection files to <number> megabytes	Indicates the maximum size for collection files. The default size is 10 MB, although you can specify a file size range from 1 MB to 99 MB. You may want to change this value to optimize the use of your backup media.
Collect files older than	Indicates the amount of time that must elapse since items were archived before they are eligible for collection.

Configuring the Backup Exec Migrator to work with a Backup Exec server

Use the following steps to configure the Backup Exec Migrator to work with a destination Backup Exec server.

Note: It is recommended that you configure two server storages devices when you configure the Backup Exec Migrator to work with Backup Exec. Configuring two storage devices lets you create a staged migration for your archived Enterprise Vault data.

See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1292.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1302.

To configure the Backup Exec Migrator to work with a Backup Exec server

- 1 At the Backup Exec server, start Backup Exec.
- 2 Create a logon account that uses the Enterprise Vault server Vault Service account credentials.

Vault Service account credentials are used so that Backup Exec and the Backup Exec Migrator can complete the migration operation.

See [“Backup Exec logon accounts”](#) on page 759.

- 3 Click the Backup Exec button and then select **Configuration and Settings**.
- 4 Click **Backup Exec Settings**, and then click **DBA-initiated Job Settings**.
- 5 Select the **DEFAULT** template, and then click **Edit**.

You can also use an existing template, or you can create a new template specifically for Enterprise Vault migrations.

- 6 Under **Storage**, select **Any disk storage** as the primary storage location for migrated data, and then set the options you want to use with the device.
- 7 Under **Migrator for Enterprise Vault**, click the down arrow next the field for **Vault Service account credentials**.
- 8 Select the logon account that you created in step 2.

See [“Migrator for Enterprise Vault options”](#) on page 1301.

- 9 In the **DBA-initiated Job Settings** dialog box, set other options as appropriate.

See [“DBA-initiated job templates”](#) on page 746.

- 10 Do one of the following:

If you want to configure staged migrations Do the following in the order listed.

- Under **Duplicate Job Settings**, check **Enable settings to duplicate backup sets for this job**.
- In the **Storage** list, select a type of storage.
- Set other options as appropriate.
See [“Duplicate job settings for DBA-initiated jobs”](#) on page 755.
- Click **OK**.

See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1292.

If you do not want to configure staged migrations Continue with step 12.

- 11 Click **OK**.
- 12 Configure Backup Exec Migrator to work with Enterprise Vault.
See “[Configuring the Backup Exec Migrator to communicate with Enterprise Vault](#)” on page 1302.

Migrator for Enterprise Vault options

The Backup Exec Migrator uses the Enterprise Vault server Vault Service account during the Backup Exec Migrator to Backup Exec server authentication process.

Table I-13 Migrator for Enterprise Vault options

Item	Description
Vault Service account credentials	Specifies the Enterprise Vault server Vault Service account credentials to use so that Backup Exec and the Backup Exec Migrator can complete the migration operation. The Vault Service account must be included in either the Administrators group or the Backup Operators group at the Backup Exec server. Note: If the Enterprise Vault server and the Backup Exec server are in different domains, a trust relationship must be established between the domains. The Vault Service account user must be a trusted user at the Backup Exec server. Trust relationships are required so that the Microsoft Security Support Provider Interface (SSPI) can authenticate the Vault Service account user. For more information on domain trust relationships, see your Microsoft documentation.
New	Lets you create a new logon account or edit an existing account. See “Backup Exec logon accounts” on page 759.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1299.

Configuring the Backup Exec Migrator to communicate with Enterprise Vault

Use the following steps to configure the Backup Exec Migrator to communicate with Enterprise Vault.

See [“Configuring the Backup Exec Migrator”](#) on page 1297.

To configure the Backup Exec Migrator to communicate with Enterprise Vault

- 1 At the Enterprise Vault server, navigate to a vault store partition from which to migrate data.
- 2 Right-click the vault store partition, and then click **Properties**.
- 3 On the **Migration** tab, check **Migrate files**.
- 4 In **Remove collection files from primary storage**, set the time period for this option to something longer than zero days.

Do not set it to zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition. More importantly, it causes the Backup Exec Migrator to create separate migration jobs for each partition folder being migrated during a migration period. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.

See [“Configuring the Backup Exec Migrator”](#) on page 1297.

- 5 Set other migration options as appropriate.
See [“Vault store partition properties - Migration options”](#) on page 1303.
- 6 On the **Advanced** tab, ensure that **Backup Exec** appears in the **List setting from** field.
- 7 In the window below the **List setting from** field, select **Backup Exec server**.
- 8 Click **Modify**.
- 9 Type the name or the IP address of the destination Backup Exec server.
- 10 Click **OK**.
- 11 Select **Backup Exec DBA-initiated template**.
- 12 Click **Modify**.

- 13 Enter the name of an existing template that uses the Enterprise Vault server Vault Service account credentials.

The template that you select must be configured to use the Enterprise Vault server Vault Service account. The template you use must also match the template name that you used when you configured the Backup Exec Migrator to work with a Backup Exec server.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1299.

- 14 Click **OK**.

- 15 Ensure that the name of the template that contains the Enterprise Vault server Vault Service account credentials appears in the **Setting** pane.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1299.

- 16 To test the communications between the Enterprise Vault server and the Backup Exec server, click **Test Configuration**.

- 17 If the test fails, ensure that you used the correct credentials for the Vault Service account, and then click **Test Configuration** again.

- 18 Click **OK** after the test successfully completes.

- 19 Click **OK**.

Vault store partition properties - Migration options

Select the Enterprise Vault migration property options that you want to use.

Table I-14 Vault store partition properties - Migration options

Item	Description
Migrate files	Lets you migrate archived Enterprise Vault data to a Backup Exec storage device. Migration can help reduce storage costs by moving collection files to tertiary storage devices. However, retrieval times can increase. See your Enterprise Vault documentation.
Migrator	Indicates the name of the migration application. Backup Exec must appear in this field.

Table I-14 Vault store partition properties - Migration options (*continued*)

Item	Description
Migrate files older than	Indicates the amount of time that must elapse since files were last modified before they are eligible for migration. See your Enterprise Vault documentation.
Remove collection files from primary storage	Indicates the age at which migrated collection files are removed from the primary storage location. Files that have been migrated to Backup Exec storage media can remain in their primary location for the period of time you specify. Note: It is recommended that you set the time period for this option to something longer than zero days, with a longer time period being best. Do not set it to zero days. Setting the time period to zero days causes the Backup Exec Migrator to create separate migration jobs in a migration period for each partition being migrated. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance. See “Configuring the Backup Exec Migrator” on page 1297.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1302.

About viewing migrated Enterprise Vault data

The Backup Exec **Backup Sets** view shows the migrated items for the Enterprise Vault partition. In the **Backup Sets** view, backup sets containing the migrated .cab files appear under a partition name that reflects the Enterprise Vault partition from which the data was migrated. Because the **Backup Sets** view displays the archived data in a read-only mode, you cannot select the data for restore. However, you can retrieve the data in the application where the data resides.

See [“Backup sets ”](#) on page 359.

Note: You can completely retrieve of all archived items that appear in the **Backup Sets** view by using the Partition Recovery Utility.

See [“About the Partition Recovery Utility”](#) on page 1306.

See [“About retrieving migrated Enterprise Vault data”](#) on page 1305.

About retrieving migrated Enterprise Vault data

All file retrieve operations start from the Enterprise Vault server console. You cannot restore archived Enterprise Vault data from Backup Exec.

When files are migrated from a partition, Enterprise Vault creates a shortcut in the partition that replaces the migrated file. The shortcut also links to the storage location of the migrated file. You retrieve files by double-clicking their shortcuts in the Enterprise Vault partition itself. If a partition retains a local copy of the migrated files, Enterprise Vault retrieves the files from the local copies. If Enterprise Vault deletes the migrated files because the partition's file retention period passes, the requested files must be retrieved from Backup Exec storage media.

Note: Backup Exec does not support restore from storage device media and from backup sets for Backup Exec Migrator for Enterprise Vault because restore of data is initiated at the Enterprise Vault server or where the data resides. Backup Sets view displays the archived data in a read-only mode; therefore, you cannot select the data for restore.

Backup Exec Migrator for Enterprise Vault does not support the restore of a resource if you back up the resource on a different storage media.

Table I-15 How migrated data is retrieved

Action	Notes
Enterprise Vault works with the Backup Exec Migrator to begin the process.	The Backup Exec Migrator identifies the Backup Exec server where the files are stored.
The Backup Exec Migrator schedules a Backup Exec restore job at the server.	Backup Exec restores the requested files.
The Backup Exec Migrator migrates the restored files to the Enterprise Vault server partition from the Backup Exec server.	The Backup Exec Migrator moves the restored files to a location specified by Enterprise Vault, using the name provided by Enterprise Vault.

The retrieval process is automatic after you start the operation at the Enterprise Vault server. It requires no user intervention other than perhaps placing a tape in the tape device if you removed the storage media.

See [“Retrieving migrated Enterprise Vault data”](#) on page 1306.

Retrieving migrated Enterprise Vault data

Use the following steps to restore migrated Enterprise Vault files.

Note: To successfully retrieve the files you want, you may need to place a tape in a tape drive at the Backup Exec server.

To retrieve migrated Enterprise Vault data

- 1 At the Enterprise Vault server, navigate to the partition where you want to retrieve the data.
- 2 Double-click the file that you want to retrieve.

About the Partition Recovery Utility

The Partition Recovery Utility is a command-line application that is automatically installed when you install the Backup Exec Agent for Windows. The utility lets you restore all of a partition's archived files from the Backup Exec storage media in a single operation. You can also use it to recover the archived partition data for each of the Enterprise Vault partitions in a disaster recovery situation.

After you use the Partition Recovery Utility, you can review recovery details by looking at the Backup Exec job history for each Enterprise Vault server where the recovery occurs.

See [“Partition Recovery Utility requirements”](#) on page 1306.

See [“Finding an archive ID”](#) on page 1307.

See [“Starting the Partition Recovery Utility”](#) on page 1307.

Partition Recovery Utility requirements

You must know the following when you use the Partition Recovery Utility:

- The vault store partition name for the data that you want to recover.
- The Archive ID of the partition data that you want to recover.
- An Enterprise Vault server user account with Vault Service Account privileges.

In addition, the Partition Recovery Utility must run at the Enterprise Vault server that originally migrated the data you want to restore.

See [“Finding an archive ID”](#) on page 1307.

See [“Starting the Partition Recovery Utility”](#) on page 1307.

Finding an archive ID

You use the archive ID of the data you want to restore along with the vault store partition name when you run the Partition Recovery Utility. The archive ID is an alpha-numeric string of considerable length.

For example, 1D69957C6D917714FB12FEA54C9A8299A1110000ev8archive.EVMBE

You can find the Archive ID listed among the properties of an archived file set.

To find an archive ID

- 1 In the left view of the the Enterprise Vault Administration Console, expand **Archives**.
- 2 Navigate the folder structure and select the folder of the type for data you want to restore.
- 3 In the right view, right-click an archive, and then select **Properties**.
- 4 On the **Advanced** tab, note the archive ID at the bottom.

See [“Starting the Partition Recovery Utility”](#) on page 1307.

Starting the Partition Recovery Utility

Use the following steps to start the Partition Recovery Utility.

To start the Partition Recovery Utility

- 1 From the Enterprise Vault server, open a Windows command prompt.
- 2 Navigate to the Enterprise Vault Agent installation directory.
For example, C:<Backup Exec install path>\Backup Exec\RAWS
- 3 Do the following:

To start the Partition Recovery Utility on all other supported Windows operating system versions	Type the following command: <pre>partitionrecovery.exe -vs <vault_store_name> -ap <archive_ID></pre>

- 4 Press **Enter**.

See [“About the Partition Recovery Utility”](#) on page 1306.

Best practices for using the Backup Exec Migrator

Consider the following best practices when you use the Backup Exec Migrator:

- It is recommended that you regularly back up the Backup Exec catalogs. In the event the catalogs become corrupt, you can restore them from backups. After you restore the catalogs, you must re-catalog the storage media on which Backup Exec Migrator data is stored. Re-cataloging the storage media ensures that the latest catalog entries are available.
- For best performance, configure the Backup Exec Migrator to migrate data to a disk-based storage, and then to a tape device by using a duplicate job. See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1292. See [“Duplicating backup sets or a job history manually”](#) on page 222.
- In the Enterprise Vault **Migration** options tab, set the time period for **Remove collection files from primary storage** to something longer than zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition.

If you set the time period to zero days, it is recommended the following:

- Increase the number of concurrent jobs that are allowed for the disk-based storage you use for migration purposes. Increase the number of concurrent jobs based on the following formula:

$$\text{<number of recommended concurrent jobs>} = \text{<number of installed tape drives plus two>}$$

For example, if you have two installed tape drives, you should configure the disk-based storage to allow four concurrent jobs.

Concurrent jobs let the Backup Exec Migrator continue to migrate data to disk storage while tape drives process duplicate jobs in a staged migration environment.

Note: You can increase the number of concurrent jobs that run by increasing the total concurrency level of the disk-based storage devices.

- It is recommended that you first collect all of the archived files in one collection and migration operation and then migrate them in the next collection and migration operation. This process helps ensure that the Backup Exec Migrator creates a single job for each migration operation, which improves the migration performance.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1287.

Troubleshooting Backup Exec Migrator and Partition Recovery Utility issues

Review the following error messages for possible solutions to errors that you may encounter:

- The Backup Exec Migrator logs migration activity in the Windows Event Viewer and in the Enterprise Vault Dtrace Utility on the Enterprise Vault server. It also logs migration activity on the Backup Exec server. The details that are provided in the log files can help you troubleshoot issues with the Backup Exec Migrator.
See [“About Backup Exec Migrator events”](#) on page 1293.
See [“About Backup Exec Migrator logs”](#) on page 1293.
- The Partition Recovery Utility cannot find any files to be recalled. There are no files to be recalled from the vault store database using the Archive ID that you provided.
- The Partition Recovery Utility operation will be terminated due to a user request.
You may have stopped the Partition Recovery Utility operation by pressing **Ctrl + C** or **Ctrl + Break**.
- The migrated file name `<file_name>` with ID `<migrated_file_id>` was not found in the Backup Exec backup sets. The recall is skipped for this file. The Partition Recovery Utility skips collection files if they already exist in the vault store database. To restore the files, delete them from the vault store database, and then run the Partition Recovery Utility again.
- The Partition Recovery Utility cannot find any partitions. Ensure that the name of the vault store is valid, and that there are partitions in the vault store. The vault store name that you provided may be invalid.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1287.

See [“About the Partition Recovery Utility”](#) on page 1306.

Enterprise Vault logon account

To back up and restore Enterprise Vault data, Backup Exec must know the user name and password for the account that is used to logon to the Enterprise Vault server. Backup Exec also uses the log on account to interact with the Enterprise Vault SQL database.

The Enterprise Vault logon account should have privileges to do the following:

- Back up and restore SQL databases.
- Communicate with the Enterprise Vault services and place Enterprise Vault into backup mode.

The logon account should also have permissions to read and write from the Enterprise Vault file system paths such as Enterprise Vault partitions and index locations. The file paths can be either fully qualified UNC paths, or the paths that are on the local drive.

Enterprise Vault options

The **Backup method** field lets you designate a backup method for when you back up Enterprise Vault data.

See [“Backup methods in Backup Exec”](#) on page 189.

Backup Exec Agent for Microsoft Active Directory

This appendix includes the following topics:

- [About the Agent for Microsoft Active Directory](#)
- [Requirements for the Agent for Microsoft Active Directory](#)
- [About backing up Active Directory and ADAM/AD LDS](#)
- [About restoring individual Active Directory and ADAM/AD LDS objects](#)

About the Agent for Microsoft Active Directory

The Backup Exec Agent for Microsoft Active Directory uses full backups for which Granular Recovery Technology (GRT) is enabled to let you restore individual Active Directory objects and attributes without performing an authoritative or non-authoritative full restore. You can also restore individual Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes.

The Agent for Microsoft Active Directory is installed as part of the Agent for Applications and Databases.

The Agent for Microsoft Active Directory works with backups of the Windows System State, where Active Directory is installed, and with ADAM/AD LDS. When you back up the Windows System State, the Active Directory is included in the backup job because Active Directory is a component of Windows System State. You can also use the Agent for Microsoft Active Directory to restore individual ADAM/AD LDS objects and attributes. If multiple ADAM/AD LDS instances are backed up, each instance appears under the Active Directory Application Mode node.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1317.

Requirements for the Agent for Microsoft Active Directory

Review the following requirements for the Agent for Microsoft Active Directory:

- The Agent for Windows must be installed on the computer on which the Active Directory is installed.
- For specific operating system requirements for the Agent for Microsoft Active Directory, refer to the Backup Exec Software Compatibility List.
- Ensure that the following option is selected on the backup job properties: **Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)**. This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups.
- To perform a GRT-enabled backup of an Active Directory Application Server on Windows 2012 R2 or later:
 - For jobs targeted to disk storage, the Backup Exec server must be running the same version or a later version of the Microsoft Windows operating system.
 - For jobs targeted to tape storage, the Hyper-V Host or equivalent host server must be running the same version or a later version of Microsoft Windows operating system.

Following are the requirements for the restore of individual objects and attributes from backup jobs that used the Agent for Microsoft Active Directory:

- You must have a full backup of Windows System State (where Active Directory is installed) or the ADAM/AD LDS.
- You must use a version of the Windows operating system that supports minifilter drivers on the Backup Exec server that runs the restore job.
- You must designate a location on the Backup Exec server where Backup Exec can temporarily place the objects and attributes that are being restored when you restore from tape.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1317.

About backing up Active Directory and ADAM/AD LDS

It is recommended that you run Active Directory and ADAM/AD LDS backups to disk storage and then copy the backups to tape. This strategy provides you with shorter backup windows and disk storage provides the most efficient method of storage for GRT-enabled backups, and the most efficient method of restore. This method also lets you administer Active Directory or ADAM/AD LDS without requiring the individual cataloging of the backed up objects and properties.

For example, if you back up to tape, you must create a temporary hard disk staging location on a local NTFS volume to restore individual items from GRT-enabled backups on tape. The data is first copied from tape to the temporary staging location before it can be restored. As such, a restore from tape takes more time. For best results, you should specifically select disk storage when you configure your GRT-enabled backup jobs.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux is installed.

Granular Recovery Technology (GRT) lets you restore individual objects and attributes from Active Directory and ADAM/AD LDS backups without performing an authoritative or non-authoritative full restore. The Granular Recovery Technology feature is enabled by default when you create a backup job. To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Server, you must use a Backup Exec server that runs Windows Server 2012 R2.

When you back up any Windows Active Directory or an ADAM/AD LDS directly to tape, objects and properties that are added or deleted during the backup will not match the individual objects and properties that are available for restore from the backup set. The backup of the database is a snapshot backup of the live Active Directory or ADAM/AD LDS database and the cataloging of the individual Active Directory or ADAM/AD LDS objects occurs after the snapshot is performed. Since the catalog operation catalogs objects and properties from the live Active Directory or ADAM/AD LDS database, object and property changes can occur after the snapshot was taken.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 205.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Microsoft Active Directory backup job options”](#) on page 1316.

See [“Granular Recovery Technology”](#) on page 739.

Editing options for Active Directory and ADAM/AD LDS backup jobs

You can edit the default settings for all Active Directory and ADAM/AD LDS backup jobs.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1317.

To edit options for Active Directory and ADAM/AD LDS backup jobs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select a backup option.
- 3 On the left, click **Microsoft Active Directory**.

- 4 Edit the following backup options for the Agent for Microsoft Active Directory as appropriate:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)

Enables the restore of individual items from full backups of the Active Directory or ADAM/ AD LDS.

This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups.

Note: You cannot restore individual objects and attributes from Active Directory backups for a read-only domain controller (RODC). You should do GRT backups and restores of the Active Directory to a writable, centralized data center domain controller.

Ensure that you meet the requirements for Granular Recovery Technology.

To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Application Server, you must use a Backup Exec server that runs Windows Server 2012 R2.

See [“Granular Recovery Technology”](#) on page 739.

Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider

Checks snapshots for data corruption. This option applies only to snapshots that are performed by the Microsoft Volume Shadow Copy Services (VSS).

Continue with backup if consistency check fails

Enables the backup job to continue even if the consistency check fails. You may want the job to continue if a backup of the database in its current state is better than no backup at all. Or you may want the job to continue if you back up a large database that may have only a small problem.

- 5 Click **OK**.

Microsoft Active Directory backup job options

You can edit the default settings for Active Directory and ADAM/AD LDS backup jobs.

Note: Only full backups of Active Directory are allowed.

See [“Editing options for Active Directory and ADAM/AD LDS backup jobs”](#) on page 1314.

Table J-1 Microsoft Active Directory backup default options

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)	Enables the restore of individual items from full backups of the Active Directory or ADAM/ AD LDS. This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups. Note: You cannot restore individual objects and attributes from Active Directory backups for a read-only domain controller (RODC). You should do GRT backups and restores of the Active Directory to a writable, centralized data center domain controller. Ensure that you meet the requirements for Granular Recovery Technology. To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Server, you must use a Backup Exec server that runs Windows Server 2012 R2. See “Granular Recovery Technology” on page 739.
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	Checks snapshots for data corruption. This option applies only to Active Directory snapshots that are performed by the Microsoft Volume Shadow Copy Service (VSS).

Table J-1 Microsoft Active Directory backup default options (*continued*)

Item	Description
Continue with backup if consistency check fails	Enables the backup job to continue even if the consistency check fails. You may want the job to continue if a backup of the Active Directory in its current state is better than no backup at all. Or you may want the job to continue if you back up a large Active Directory that may have only a small problem.

See [“About backing up Active Directory and ADAM/AD LDS”](#) on page 1313.

About restoring individual Active Directory and ADAM/AD LDS objects

When you restore Active Directory and ADAM/AD LDS objects from tape, you must specify an on-disk staging location where the objects will be placed prior to being restored. The staging location must be a path on a local NTFS volume on the Backup Exec server running the restore job and the Backup Exec service account must also have access to it.

By default, the Agent for Microsoft Active Directory restores deleted Active Directory or ADAM/AD LDS objects from the Active Directory Deleted Objects container if their tombstone lifetimes have not passed.

When objects in Active Directory are deleted, they are removed from their current Active Directory or ADAM/AD LDS container, converted into tombstones, and then placed in the Active Directory Deleted Objects container where their tombstone lifetime is monitored. After their tombstone lifetime passes, the tombstones are purged from the Active Directory Deleted Objects container, which permanently deletes the objects from the Active Directory and ADAM/AD LDS databases.

The Agent for Microsoft Active Directory lets you restore tombstoned objects from the Active Directory Deleted Objects container in the following situations:

- Their tombstone lifetimes have not passed.
- They have not been purged from the Deleted Objects container.

When you restore Active Directory user objects, you must reset the object’s user password and then re-enable the object’s user account. For ADAM/AD LDS user objects, you must reset the object’s user password and then re-enable the object’s

user account. For Active Directory user objects, use the Microsoft Active Directory Users and Computers application. For ADAM/AD LDS user objects, use ADSI Edit.

For Active Directory computer objects, you must reset the object's account.

Some objects in the Active Directory Configuration Partition node cannot be reanimated from the Active Directory Deleted Objects container. However, recreated objects may not be recognized by some applications.

Note: When you restore ADAM/AD LDS data, Backup Exec stops the ADAM/AD LDS instance you want to restore before the restore job starts. However, Backup Exec does not restart the ADAM/AD LDS instance when the restore job completes because post-processing jobs, such as authoritative restores using Adamutil.exe, may be needed. You must restart the ADAM/AD LDS instance. If Backup Exec cannot stop the ADAM/AD LDS instance or if Backup Exec cannot restore all of the ADAM/AD LDS data, the restore fails.

For more information, see your Microsoft Active Directory documentation.

Resetting the Active Directory computer object and the computer object account

In Active Directory, computer objects are derived from user objects. Some attributes that are associated with a computer object cannot be restored when you restore a deleted computer object. The attributes can only be restored if the attributes were saved through schema changes before the computer object was originally deleted. Because computer object credentials change every 30 days, the credentials from the backup may not match the credentials that are stored on the actual computer.

Note: To reset a computer object, you must use the Microsoft Active Directory Users and Computers application.

For more information on resetting a computer object, see your Microsoft Active Directory Users and Computers application documentation.

If a computer object's **userAccountControl** attribute was not preserved before the object was deleted, you must reset the object's account after you restore the object.

To reset the Active Directory computer object account

- 1 Remove the computer from the domain.
- 2 Re-join the computer to the domain. The SID for the computer remains the same since it is preserved when you delete a computer object. However, if the object's tombstone expires and a new computer object is recreated, the SID is different.

Recreating purged Active Directory and ADAM/AD LDS objects

You can attempt to recreate deleted Active Directory objects and ADAM/LDS objects after they have been purged from the **Active Directory Deleted Objects** container by restoring the object from a previous Active Directory backup.

You can attempt to recreate the deleted objects if their tombstone lifetimes have passed and the objects have been purged from the Active Directory Deleted Objects container.

However, you should be aware of the following:

- Most applications will not recognize a recreated object since recreated objects are not identical to the original deleted object. Recreated objects are assigned new global unique identifiers (GUIDs) and security identifiers (SIDs) that cannot be identified by the applications that created the original object.
- Attributes created by the Windows operating system cannot be recreated when a purged object is recreated. Hence, objects that rely on attributes set by the operating system will not be recognized by Windows when the objects are recreated.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1317.

Backup Exec Central Admin Server Feature

This appendix includes the following topics:

- [About the Central Admin Server feature](#)
- [Requirements for installing CAS](#)
- [How to choose the location for CAS storage and media data](#)
- [About installing the Central Admin Server feature](#)
- [Upgrading an existing CAS installation](#)
- [Changing a Backup Exec server to a central administration server](#)
- [Changing a Backup Exec server to a managed Backup Exec server](#)
- [Deleting a managed Backup Exec server from a CAS environment](#)
- [Renaming a central administration server](#)
- [Renaming a managed Backup Exec server](#)
- [How to reduce network traffic in CAS](#)
- [CAS distributed, centralized, and replicated catalog locations](#)
- [Changing the settings for a managed Backup Exec server](#)
- [What happens when CAS communication thresholds are reached](#)
- [Alerts and notifications in CAS](#)
- [Enabling managed Backup Exec servers to use any available network interface card](#)

- [About job delegation in CAS](#)
- [About adding storage devices in a CAS environment](#)
- [How data lifecycle management \(DLM\) works in a CAS environment](#)
- [Obtaining media audit information for a managed Backup Exec server](#)
- [How to use Backup Exec server pools in CAS](#)
- [How centralized restore works in CAS](#)
- [About recovering failed jobs in CAS](#)
- [Pausing or resuming a managed Backup Exec server](#)
- [Stopping or starting Backup Exec services for a managed Backup Exec server](#)
- [Viewing managed Backup Exec server properties](#)
- [Viewing the settings for a central administration server](#)
- [Disaster recovery in CAS](#)
- [Troubleshooting CAS](#)
- [Running the Backup Exec Utility for CAS operations](#)
- [Uninstalling Backup Exec from the central administration server](#)
- [Uninstalling Backup Exec from a managed Backup Exec server](#)

About the Central Admin Server feature

The Backup Exec Central Admin Server feature (CAS) enables a central administration server to delegate jobs to managed Backup Exec servers across the network. Job delegation is the automatic load balancing of jobs across available managed Backup Exec servers in the CAS environment. If your organization includes more than one Backup Exec server, you can benefit from using CAS. Refer to the Price and Licensing Guide for more information about which Backup Exec editions allow using the CAS feature.


All backup information in the CAS environment can be centralized on the central administration server. The managed Backup Exec servers perform the actual processing of backup and restore jobs. You create jobs on the central administration server and then delegate the jobs to run on a managed Backup Exec server. The jobs are delegated, or load-balanced, across the available storage devices on the managed Backup Exec server. Multiple Backup Exec servers can

share a storage device when the sharing option is enabled. Centralized restore jobs can also be delegated to managed Backup Exec servers. Additionally, the central administration server can function as a managed Backup Exec server and process delegated jobs. A managed Backup Exec server can also run the jobs that are created locally at its local administration console.

For information about the best practices to use Backup Exec Central Admin Server Option (CASO), refer to *Backup Exec Best Practices*.


The following graphic shows a local area network (LAN) environment with a central administration server and several managed Backup Exec servers.

Figure K-1 CAS-configured Backup Exec environment - LAN



The same communications that occur over a LAN between the central administration server and the managed Backup Exec servers take place over a WAN.

Figure K-2 CAS-configured Backup Exec environment - WAN



See [“Sharing storage devices”](#) on page 557.

See [“How to choose the location for CAS storage and media data ”](#) on page 1324.

See [“About installing the Central Admin Server feature”](#) on page 1326.

See [“Upgrading an existing CAS installation”](#) on page 1335.

Requirements for installing CAS

The system requirements for the Central Admin Server feature (CAS) are the same as the minimum requirements for Backup Exec, with the exception of RAM. However, processor speed, memory, and disk space requirements may increase based on the number of managed Backup Exec servers, the number of servers being backed up, and the amount of catalog storage space that is required.

On the computer on which you install the central administration server, 1 GB RAM is required, although 2 GB RAM or more is recommended for better performance. Other applications on the Backup Exec server also require a certain amount of physical RAM to function properly. The requirements for RAM may

also increase when the central administration server manages more Backup Exec servers or tape hardware.

Before you install CAS, do the following:

- Ensure that you have administrative rights on computers on which you want to install CAS.
- Ensure that when you install CAS on Backup Exec servers in multiple domains, the Backup Exec service account is in the trusted domain, and has administrative rights on all of the Backup Exec servers that you want to use as managed Backup Exec servers.
If the Backup Exec Database for the central administration server is installed on a SQL Server instance on a different computer, the account must be a domain account with local administrative privileges on that computer as well.
- Ensure that the central administration server and the managed Backup Exec servers are part of a domain or domains. CAS is not supported in a workgroup.
- Use only NetBIOS computer names for managed Backup Exec servers and central administration servers. You cannot enter fully qualified domain names or IP addresses as server names.
- Ensure that you have the appropriate licenses for Backup Exec. A license for Backup Exec is required in addition to a license for CAS.
- Ensure that when you install a managed Backup Exec server, both the logged-in user and the Backup Exec service account for the managed Backup Exec server have administrative rights on the central administration server.

See [“System requirements for Backup Exec”](#) on page 45.

See [“About installing the Central Admin Server feature”](#) on page 1326.

How to choose the location for CAS storage and media data

During the installation of the managed Backup Exec server feature, you can choose the location of the managed Backup Exec server’s storage and media data.

The following table compares how CAS tasks are performed depending on the location of the managed Backup Exec server’s storage and media data:

Table K-1 Comparison of CAS tasks

Task	Storage and media data on the central administration server	Storage and media data on the managed Backup Exec server
Delegate jobs from the central administration server to the managed Backup Exec server	Yes	No. Instead, you can create jobs on the central administration server, and then copy them to the managed Backup Exec server.
Manage storage devices and media on the managed Backup Exec server from the central administration server	Yes	No
Hold, delete, run, cancel, and change the priority of copied jobs from the central administration server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Monitor the jobs that are created on the local managed Backup Exec server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Send job status updates, job logs, and job histories to the central administration server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Centralize, distribute, or replicate the catalog	Yes	No Only a distributed catalog location can be selected.

Table K-1 Comparison of CAS tasks (continued)

Task	Storage and media data on the central administration server	Storage and media data on the managed Backup Exec server
Run centralized restore	Yes	Yes You can browse the backup sets and run restore operations for the managed Backup Exec server from the central administration server.

Note: In a CAS environment, you can add an NDMP Server only to a central administration server or a managed Backup Exec server on which the storage and media database is located.

- See [“Upgrading an existing CAS installation”](#) on page 1335.
- See [“About the Central Admin Server feature”](#) on page 1321.
- See [“Running the Backup Exec Utility for CAS operations”](#) on page 1380.

About installing the Central Admin Server feature

The Central Admin Server feature is installed as part of the Enterprise Server feature. After you enter a license for Backup Exec, on the **Configure features** panel, you must expand the **Backup Exec Features** item, and then expand the **Enterprise Server Feature** item to select the Central Admin Server feature for installation. When you select the Central Admin Server feature for installation, the central administration server is installed. After the central administration server is installed, you can install managed Backup Exec servers.

Note: You must use the custom installation option in the installation wizard to install CAS. The typical installation option does not support the installation of CAS.

- See [“Push-installing a managed Backup Exec server from the central administration server ”](#) on page 1327.
- Before you start the installation, review the information about the location of storage and media data.

See [“How to choose the location for CAS storage and media data ”](#) on page 1324.

A managed Backup Exec server may be installed outside the firewall that the central administration server is installed in or in a different firewall. However, you must set up a static port for SQL Server and create an alias for the managed Backup Exec server.

See [“Installing a managed Backup Exec server across a firewall ”](#) on page 1333.

Push-installing a managed Backup Exec server from the central administration server

After you install the central administration server, you can push-install the managed Backup Exec server feature to a standalone server.

Before you install a managed Backup Exec server, decide where to locate the storage and media database for it. During the installation of the managed Backup Exec server, you can choose the location of the managed Backup Exec server's storage and media data. Your choice affects how you can manage jobs in the CAS environment.

See [“How to choose the location for CAS storage and media data ”](#) on page 1324.

To push-install a managed Backup Exec server from the central administration server

- 1 From the central administration server, click the Backup Exec button, and then select **Installation and Licensing**.
- 2 Select **Install Agents and Backup Exec Servers on Other Servers**.
- 3 In the installation wizard, click **Add**, and then select either **Add a Single Computer** or **Add Multiple Computers with the Same Settings**.
- 4 Select **Backup Exec**, and then click **Next**.
- 5 In the **Remote computer** field, type the name of the managed Backup Exec server that you want to add, or click **Browse Remote Computers** to locate the server.
- 6 Click **Add to List**.

This option is not necessary if you selected **Add a Single Computer** in step 3.

7 Under **Remote computer credentials** , complete the fields as follows:

User Name	Type the user name for an account that has administrative rights on the remote computer.
Password	Type the password for an account that has administrative rights on the remote computer.
Domain	Select the domain in which the remote computer is located.

8 Click **Next**.

9 Select one of the following methods to enter license keys:

To enter entitlement IDs manually	In the Enter an Entitlement ID field, type an entitlement ID from your sales certificate, and then click Add .
To import a license file	Click Import License File , and then navigate to the location of your .slf file.
To install a trial version	Do not enter entitlement IDs or import license files. Proceed to the next step.

10 Click **Next**.

11 After your entitlement IDs are validated, click **Next**.

12 On the list of features to install, expand **Backup Exec**, and then select **Managed Backup Exec server**.

13 Do one of the following:

To change the directory where the Backup Exec files are installed	In the Destination Folder field, type the name of the directory.
To accept the default directory (recommended)	Proceed to the next step.

It is recommended that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec is uninstalled.

14 Click **Next**.

15 Provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

- 16 On the **Choose SQL Server** panel, choose the location to store the Backup Exec Database, and then click **Next**.
- 17 In the **Central Administration Server** field, type the name of the central administration server that will manage this managed Backup Exec server.
Use only NetBIOS computer names for managed Backup Exec servers and central administration servers. You cannot enter fully qualified domain names or IP addresses as server names.
- 18 Select from the following options to determine how storage devices and data are managed:

See [“Managed Backup Exec Server Configuration options”](#) on page 1331.

Centrally managed Backup Exec server Select this option to enable the central administration server to manage this Backup Exec server, its storage devices, media, and job delegation. This option also enables this Backup Exec server to share storage devices with other managed Backup Exec servers.

Unrestricted access to catalogs and backup sets for restore Select this option to enable this managed Backup Exec server to have unrestricted access to all centrally stored catalogs. This option also enables this managed Backup Exec server to restore data from any backup set on any storage devices that it shares.

This option can be selected only if the **Centrally managed Backup Exec server** option is selected. Selecting both of these options enables the central administration server to have the greatest amount of control over this managed Backup Exec server.

Locally managed Backup Exec server Select this option to enable the central administration server to monitor this managed Backup Exec server and create restore jobs for it. However, the server and its devices, media, and backup jobs are controlled locally.

- 19 Click **Next**.
- 20 Select the device drivers that you want to use, and then click **Next**.

21 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer	Click Add , and then click Add a Single Server .
To manually add multiple remote computers	Click Add , and then click Add Multiple Servers with the Same Settings .
To add multiple remote computers by importing an existing list of computers	Click Import and Export, and then select one of the following options: ■ Select Import from File to enable Backup Exec to add the names of the remote computers from a selected list. ■ Select Import Servers Published to this Backup Exec Server to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server. You must enter remote computer logon credentials for the list of remote computers.
To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	Verify that Save the server list for future remote install sessions is checked. This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers.
To save this list of remote computers to an XML file	Click Import and Export, and then click Export to File. You can select the location to save the XML file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.

To fix the errors that were located during the validation Right-click the name of the computer, and then click **Fix Error**.

To enable Backup Exec to attempt to re-validate an invalid remote computer Right-click the name of the computer, and then click **Retry Validation**.

- 22 After all of the computers are validated, click **Next**.
- 23 Read the Backup Exec installation review, and then click **Install**.
- 24 Click **Next**, and then click **Finish**.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Managed Backup Exec Server Configuration options

The following information can help you to determine which configuration options to choose when you install a managed Backup Exec server.

The following information applies if you select the option **Centrally managed Backup Exec server** and also select the option **Unrestricted access to catalogs and backup sets for restore**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is required between the managed Backup Exec server and the central administration server.
- The catalogs are centralized and are stored on the central administration server. Note that this combination of options may not be suitable if you have a low-bandwidth network connection to the central administration server.
- This managed Backup Exec server can access and restore backup sets for all storage devices that it shares with other Backup Exec servers.

Note: A managed Backup Exec server that has been configured with the catalog mode as 'centralized unrestricted' can only browse its own sets. This server cannot browse the backup sets created by any other managed Backup Exec server or the central administration server.

- The backup jobs that are created on the central administration server can be load-balanced and delegated to this managed Backup Exec server.
- A rolling upgrade cannot be performed with this configuration. This managed Backup Exec server must be upgraded at the same time as the central administration server.

The following information applies if you select the option **Centrally managed Backup Exec server**, but do not select the option **Unrestricted access to catalogs and backup sets for restore**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is required between the managed Backup Exec server and the central administration server.
- The catalogs are in distributed mode by default, but can be changed. The catalogs for the jobs that run on this managed Backup Exec server are stored locally.
- This managed Backup Exec server can access and restore any backup set that is stored on the storage devices that it hosts, regardless of which Backup Exec server ran the backup job. However, for the shared storage devices that other Backup Exec servers host, this managed Backup Exec server can access and restore only the backup sets that were created from backup jobs it ran.
- The backup jobs that are created on the central administration server can be load-balanced and delegated to this managed Backup Exec server.
- This option is recommended for use with the private cloud configuration.

The following information applies if you select the option **Locally managed Backup Exec server**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is not required between the managed Backup Exec server and the central administration server. Therefore, this option may be useful when a very low-bandwidth connection exists between the managed Backup Exec server and the central administration server. It may also be useful when the managed Backup Exec server cannot always be connected to the central administration server.
- The catalogs are in distributed mode by default. The catalogs for the jobs that run on this managed Backup Exec server are stored locally.
- The central administration server does not delegate jobs to this managed Backup Exec server.
- This managed Backup Exec server cannot be used in a private cloud configuration.

See [“Push-installing a managed Backup Exec server from the central administration server”](#) on page 1327.

Installing a managed Backup Exec server across a firewall

A managed Backup Exec server may be installed outside the firewall that the central administration server is installed in or in a different firewall.

The following rules apply to the managed Backup Exec servers that are installed across a firewall:

- Port 3527 must be open in both directions to enable communication for the Backup Exec Server service.
- Port 10000 must be open for the Agent for Windows, which allows browsing for remote selections.
- A SQL port must be open in both directions to the central administration server's database to enable database connections.
- A static port must be used.

The Backup Exec SQL instance is configured by default to use a dynamic port. Each time SQL Server is started, the port number can change. You must change the dynamic port to a static port. After you change the configuration of the port from dynamic to static, you must add the static port to the Windows Firewall Exceptions list.

See your Windows operating system documentation.

See ["To change the dynamic port for a SQL Express instance to a static port and create an alias for the managed Backup Exec server"](#) on page 1333.

See ["To open a SQL port and create an alias for a managed Backup Exec server"](#) on page 1334.

To change the dynamic port for a SQL Express instance to a static port and create an alias for the managed Backup Exec server

- 1 On the central administration server, click **Start > All Programs > Microsoft SQL Server > Configuration Tools > SQL Server Configuration Manager**.
- 2 Expand **SQL Server Network Configuration**.
- 3 Click **Protocols for BKUPEXEC**, and then in the right pane, double-click **TPC/IP**.
- 4 On the **TCP/IP Properties** dialog box, click the **IP Addresses** tab.
- 5 Under **IPAll**, in **TCP Dynamic Ports**, remove the value and leave the field blank.
- 6 Under **IPAll**, type in a port number in **TCP Port**.

The port number can be between 1025 and 65535 and must not be in use by another application.

- 7 Under the heading for the specific network interface card that is being used, such as IP1 or IP2, change **Enabled** from **No** to **Yes**.
- 8 Under that same heading, in **TCP Dynamic Ports**, remove the value of 0, and type the same port number you entered for **TCP Port**.
- 9 Click **Apply**.
- 10 Restart the Backup Exec and SQL services.
- 11 On the managed Backup Exec server, click **Start > All Programs > Microsoft SQL Server > Configuration Tools > SQL Server Configuration Manager**.
- 12 Expand **SQL Native Client Configuration**.
- 13 Click **Aliases**, and then double-click the alias name that contains the central administration server name and the Backup Exec SQL instance name.
- 14 On the alias properties dialog box, enter the appropriate information as described in the following table:

Alias Name	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.
Port No	Type the port number of the remote Backup Exec SQL Server instance that you noted in the previous procedure.
Protocol	Select TCP/IP .
Server	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.

- 15 Click **Apply**, and then click **OK**.
- 16 Close the SQL Server Configuration Manager utility.

To open a SQL port and create an alias for a managed Backup Exec server

- 1 On the central administration server, click **Start > All Programs > Microsoft SQL Server > Configuration Tools > SQL Server Configuration Manager**.
- 2 Expand **SQL Server Network Configuration**, and then click **Protocols** for the SQL Server instance that is in use for the central administration server.
- 3 Under **Enabled Protocols**, select **TCP/IP**, and then click **Properties**.
- 4 Double-click **TCP/IP**, and then click the **IP Addresses** tab.
- 5 Write down the TCP Dynamic port number.

- 6 On the managed Backup Exec server, to create an alias for the managed Backup Exec server, go to `\Windows\System32` and double-click **cliconfg.exe**.
- 7 On the **Alias** tab, click **Add**.
- 8 In the **Server alias** field, type: *server name\instance name*
- 9 Under **Network libraries**, select **TCP/IP**.
- 10 In the **Server name** field, type: *server name\instance name*
- 11 Uncheck **Dynamically determine port**.
- 12 In the **Port number** field, type the port number of the remote Backup Exec SQL Server instance.

Upgrading an existing CAS installation

In an existing CAS environment, upgrade the central administration server, and then upgrade the managed Backup Exec servers.

If necessary, you can perform rolling upgrades in the CAS environment. A rolling upgrade lets you upgrade the central administration server from the previous version to the current version first, and then upgrade the managed Backup Exec servers from the previous version to the current version over a period of time. You must have the most recent Backup Exec feature pack to perform rolling upgrades.

Note: Forward compatibility is not supported in rolling upgrades. Therefore, any system that runs the previous version cannot protect a system that runs the current version.

It is recommended that you do not keep a mix of versions in the CAS installation for an extended time. Key functionality for administering managed Backup Exec servers is missing in a mixed-version environment, which decreases your ability to properly administer the CAS environment.

Note: Do not rename any managed Backup Exec servers or the central administration server during the upgrade process. Although you can rename the servers before you upgrade, it is recommended that you rename them after the upgrade process has completed.

After you upgrade the central administration server to the current version, the following operations are supported on managed Backup Exec servers that run the previous version:

- Backup
- Restore
- Inventory
- Catalog

To upgrade an existing CAS installation

- 1 Verify that the latest feature pack for Backup Exec is installed.
- 2 Place all scheduled jobs on hold on the central administration server and the managed Backup Exec servers.
See [“Holding jobs”](#) on page 266.
- 3 Allow all active jobs to complete.
- 4 Run a database maintenance job to delete job histories and catalogs that you no longer need in order to shorten the upgrade window.
- 5 Stop all Backup Exec services on each managed Backup Exec server
- 6 From the installation media browser, select the option to install Backup Exec.
- 7 On the **Welcome** panel, click **Next**.
- 8 Select **I accept the terms of the license agreement**, and then click **Next**.
- 9 Check **Local Install**, and then click **Install Backup Exec software and features**.
- 10 Click **Next**.
- 11 Follow the prompts in the wizard.
- 12 On the **Back Up Existing Catalog and Data** page, enter or browse to a directory to which all existing catalogs and data will be backed up. The default location is:

`C:<Backup Exec install path>\Backup Exec\Data`

If you do not want to keep previous catalogs and data, click **Do not back up previous data and catalogs**.

- 13 Click **Next** to continue.

An upgrade summary is displayed. When the upgrade is complete, communication with the managed Backup Exec servers is automatically enabled.

- 14 Release the hold on all the jobs.
See [“Removing the hold on jobs”](#) on page 267.
- 15 Upgrade some or all of the managed Backup Exec servers.

Note: You should enable the **Allow only Kerberos authentication from the Remote Administration Console** check box only after all the Managed Backup Exec Server (MBES) and CAS servers are upgraded to the same version.

See [“CAS distributed, centralized, and replicated catalog locations”](#) on page 1344.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1346.

Changing a Backup Exec server to a central administration server

You can change a standalone Backup Exec server to a central administration server.

To change a Backup Exec server to a central administration server

- 1 On the Backup Exec server that you want to be the central administration server, start Backup Exec.
- 2 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install features and Licenses on this Backup Exec Server**.
- 3 Select one of the following methods to enter licenses:

To enter licenses manually Do the following in the order listed:

- Type an entitlement ID into the **Enter an Entitlement ID** field.
- Click **Add to list**.
- Repeat for each entitlement ID for each feature or agent that you want to add.

To import licenses from a file Do the following in the order listed:

- Click **Import License File**.
- Select the license file.

To install a trial version

Proceed to the next step.

A license key is not required for a fully functional Trial version.

- 4 Click **Next**.
- 5 On the features list, expand **Backup Exec features**, expand **Enterprise Server feature**, and then select **Central Admin Server feature**.
- 6 Click **Next**.
- 7 Read the Backup Exec installation review, and then click **Install**.
- 8 Click **Finish**.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1338.

Changing a Backup Exec server to a managed Backup Exec server

To change a Backup Exec server to a managed Backup Exec server, you set the central administration server that will manage the Backup Exec server.

If the managed Backup Exec server does not appear on the **Storage** tab after you follow these instructions, and if your network contains firewalls, you may need to open some ports between the central administration server and the managed Backup Exec server.

To change a Backup Exec server to a managed Backup Exec server

- 1 At the standalone Backup Exec server, click the Backup Exec button, select **Installation and Licensing**, and then select **Install features and Licenses on this Backup Exec Server**.
- 2 On the **Add Licenses** panel, click **Next**.
- 3 On the **Configure features** panel, under **Backup Exec**, select **Managed Backup Exec server**, and then click **Next**.
- 4 When you are prompted to select an additional language, click **Next**.
- 5 Enter the name of the central administration server that you want to manage this Backup Exec server.

- 6 Under **Managed Backup Exec Server Configuration**, select the appropriate option, and then click **Next**.

Centrally managed Backup Exec server Select this option to enable the central administration server to manage this Backup Exec server, its storage devices, media, and job delegation. This option also enables this Backup Exec server to share storage devices with other managed Backup Exec servers.

Unrestricted access to catalogs and backup sets for restore Select this option to enable this managed Backup Exec server to have unrestricted access to all centrally stored catalogs. This option also enables this managed Backup Exec server to restore data from any backup set on any storage devices that it shares.

This option can be selected only if the **Centrally managed Backup Exec server** option is selected. Selecting both of these options enables the central administration server to have the greatest amount of control over this managed Backup Exec server.

Locally managed Backup Exec server Select this option to enable the central administration server to monitor this managed Backup Exec server and create restore jobs for it. However, the server and its devices, media, and backup jobs are controlled locally.

- 7 Click **Install**.
- 8 After the installation completes, click **Finish**.

See [“Changing a Backup Exec server to a central administration server”](#) on page 1337.

Deleting a managed Backup Exec server from a CAS environment

If communications are still active between the managed Backup Exec server and the central administration server, you can remove a managed Backup Exec server

by deleting it from the **Storage** tab. When you delete a managed Backup Exec server from the **Storage** tab, the server is changed to a standalone Backup Exec server. After the managed Backup Exec server is removed from the CAS environment, it still remains on the list of servers on the **Backup and Restore** tab, so it can still be backed up and restored as a standalone server.

Note: If the central administration server experiences a disaster, communications are lost between the managed Backup Exec server and the central administration server and the managed Backup Exec server cannot be deleted from the **Storage** tab. However, you can change the managed Backup Exec server to a standalone Backup Exec server by using the Windows Change Program feature.

See [“Disaster recovery in CAS”](#) on page 1377.

To delete a managed Backup Exec server from a CAS environment

- 1 If the storage devices on the managed Backup Exec server that you want to delete are shared with other managed Backup Exec servers, on the **Storage** tab, right-click the device, and then click **Share**. Uncheck the servers that share the device.

You must remove the sharing reference to avoid overwriting the media.

Note: If any devices are shared by Fibre Channel SAN, disconnect the devices that are on the managed Backup Exec server before you delete it.

- 2 On the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to delete from the CAS environment and convert to a standalone Backup Exec server.
- 3 Select **Delete**.

- 4 Click **Yes** to confirm that you want to convert the server to a standalone Backup Exec server.
- 5 After you receive an alert on the central administration server that confirms the server was deleted, restart the server that was changed to a standalone Backup Exec server.

Note: If Backup Exec is unable to complete the request, you can choose to remove the reference to the managed Backup Exec server from the central administration server's database. Then, when the managed Backup Exec server is brought online again, Backup Exec automatically adds the managed Backup Exec server back to the central administration server's database. At that time, you can perform this procedure again.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1338.

Renaming a central administration server

If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. You should select the icon with the old server name to restore any data that you backed up before you changed the server name. You should select the icon with the new server name to restore any data that you backed up after you changed the server name.

After you rename a central administration server, that catalogs folder uses the new central administration server's name. An automatic process called catalog self-healing modifies the catalog files on the server and the catalog metadata in the Backup Exec Database.

Note: If you want to rename a central administration server that does not have any managed Backup Exec servers associated with it, you do not need to follow this procedure. Instead, you can use the standard Windows renaming process.

To rename a central administration server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server, and then click **Delete**.

The managed Backup Exec server is not deleted from Backup Exec. It is removed only from the CAS environment, so it becomes a standalone Backup Exec server.
- 2 Repeat step 1 for every managed Backup Exec server that the central administration server manages.
- 3 Rename the server by using the standard Windows renaming process.
- 4 To rejoin the managed Backup Exec servers to the central administration server, install the managed Backup Exec server feature on the servers that were deleted from the **Storage** tab in steps 1 and 2.

Note: Restart the server when the installation is complete.

- 5 Verify that the managed Backup Exec servers are rejoined to the renamed central administration server.

Renaming a managed Backup Exec server

Before you can rename a managed Backup Exec server, you must delete it from the CAS environment so that it becomes a standalone Backup Exec server. After you rename the server, you must reinstall the managed Backup Exec feature to the server, and then redelegate all of the jobs to the new managed Backup Exec server name.

If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. You should select the icon with the old server name to restore any data that you backed up before you changed the server name. You should select the icon with the new server name to restore any data that you backed up after you changed the server name.

After you rename a managed Backup Exec server, the catalogs folder uses the new managed Backup Exec server's name. An automatic process called catalog self-healing modifies the catalog files on the server and the catalog metadata in the Backup Exec Database.

To rename a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to rename.

- 2 Click **Delete** to remove the managed Backup Exec server from the CAS environment.

The managed Backup Exec server is not deleted from Backup Exec. It is removed only from the CAS environment, so it becomes a standalone Backup Exec server. The server must be removed from the CAS environment before it can be renamed.

- 3 Restart the services on the server that you want to rename.
- 4 Rename the server by using the standard Windows renaming process.
- 5 Restart the server.

After the server is restarted, Backup Exec may prompt you to restart the Backup Exec services and the Backup Exec deduplication services. You must choose the new server name on the **Backup Exec Services Manager** dialog box.

- 6 To rejoin the server to the central administration server, install the managed Backup Exec server feature on the renamed server.

Note: Restart the server when the installation is complete.

- 7 Verify that the renamed managed Backup Exec server is rejoined to the central administration server.
- 8 Delegate all jobs that were associated with the managed Backup Exec server prior to the renaming to the renamed managed Backup Exec server.

How to reduce network traffic in CAS

To accommodate a low-bandwidth network connection or to reduce network traffic, you can do the following:

- Reduce the frequency of job status updates that are sent from the managed Backup Exec servers to the central administration server.
- Prevent the central administration server from monitoring the jobs that are created on the local managed Backup Exec servers.
- Reduce the frequency that job logs and job histories are sent from the managed Backup Exec servers to the central administration server.

- Increase the amount of time that Backup Exec waits before changing the Backup Exec server's status if the Backup Exec server becomes unresponsive.
- Keep the catalogs on the managed Backup Exec server (distributed). If there is a persistent network connection between the central administration server and the managed Backup Exec server, then you can browse the catalog and perform restore operations from both servers, regardless of the catalog location.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1346.

CAS distributed, centralized, and replicated catalog locations

In the CAS environment, you can choose the catalog location. Regardless of the catalog location, if a persistent network connection is available between the central administration server and the managed Backup Exec server, then you can browse the backup sets in the catalog and perform restore operations from both servers.

The following catalog locations are available:

Table K-2

Item	Description
Distributed	Image files, which are small files that contain information about the backup set, are distributed to the central administration server from every managed Backup Exec server. History files, which contain detailed information about the backup set, remain on the managed Backup Exec server. Note: It is important that you back up the catalog files on the managed Backup Exec server since most catalog information is kept here when the distributed catalog location is used. When the catalog is distributed, the view of the restore selections on the central administration server displays only the backup set at the volume level. Backup set details are not displayed if the managed Backup Exec server that created this backup set is not available, but the whole volume can be restored from the central administration server. A distributed catalog provides increased performance, default centralized restore capability, and decreased network traffic. If a managed Backup Exec server does not have a persistent connection to the central administration server, then whenever the managed Backup Exec server does connect, the image files in the catalog are automatically distributed to the central administration server. The temporary increase in network traffic that is caused by the catalog distribution is not significant.
Centralized	All catalog files and information for the managed Backup Exec server are kept on the central administration server.

Table K-2 (continued)

Item	Description
Replicated	All catalog files are replicated from the managed Backup Exec server to the central administration server. Both the managed Backup Exec server and the central administration server store the catalogs that are produced by the managed Backup Exec server. Deletions of catalog files are replicated between the managed Backup Exec server and the central administration server only when the catalog files are deleted by Backup Exec according to the catalog settings. If catalog files on the managed Backup Exec server are deleted as a result of a backup job or a manual deletion, the deletions are replicated the next time that the catalogs are synchronized.

When choosing the catalog location, consider the following:

- If there is enough available disk space on the managed Backup Exec server to keep a distributed or replicated catalog.
- If there is enough network bandwidth to handle the traffic that is generated by a centralized or replicated catalog. Centralized and replicated catalogs require a high-bandwidth network connection.
- If it is important for your data recovery needs to keep catalog information in one location. For example, when the catalog location is centralized or replicated, all catalog information is kept in one location, which makes it easier to back up. When the catalog location is distributed, most catalog information is kept on the managed Backup Exec server.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1346.

Changing the settings for a managed Backup Exec server

The settings for a managed Backup Exec server determine how the managed Backup Exec server communicates and interacts with the central administration server. For example, you can change the connection type, the catalog location,

and the job reporting and monitoring functionality. You can change the settings for a managed Backup Exec server at any time.

Note: You may need to restart the services on the managed Backup Exec server after you change the settings. For example, if you change the catalog location, you must restart the services to enable the change to take effect.

To change the settings for a managed Backup Exec server

- 1 On the central administration server, on the **Storage** tab, double-click the managed Backup Exec server.
- 2 In the left pane, select **Settings**.
- 3 Select the appropriate options.

Connection settings

Select one of the following types of connections with the central administration server:

- **Fast connection**
Configures frequent communications between the central administration server and the managed Backup Exec server. By default, when you choose this setting, job status updates are sent every 10 seconds to the central administration server. Job logs and job histories are sent whenever a job on the managed Backup Exec server completes.
- **Slow connection**
Configures less frequent communications between the central administration server and the managed Backup Exec server. By default, when you choose this setting, job status updates are sent every 120 seconds to the central administration server. Job logs and job histories are sent only when a job on the managed Backup Exec server fails.
- **Custom**
Lets you change the thresholds that trigger the communication statuses when managed Backup Exec servers become unresponsive. You can also set how often the managed Backup Exec server sends active job status updates to the central administration server. The frequency affects the network traffic.

Communication stalled

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Indicate the amount of time to allow before the managed Backup Exec server's status changes to Communication Stalled if the managed Backup Exec server is unresponsive.

The central administration server does not delegate jobs to the managed Backup Exec server when it has a status of Communication Stalled. Job delegation resumes if the managed Backup Exec server returns to an Enabled status before the threshold is exceeded.

The default threshold is five minutes.

No communication

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Indicate the amount of time to allow before the managed Backup Exec server's status changes from Communication Stalled to No Communication.

When the status of the managed Backup Exec server changes from Communication Stalled to No Communication, the central administration server marks the active jobs on the managed Backup Exec server as Failed. The custom error-handling rule Recovered Jobs is applied to any job that is active when the No Communication status appears.

The default threshold is 15 minutes.

Send active job status updates to the central administration server

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Select this option to enable the managed Backup Exec server to send a job status update to the central administration server. You can adjust the number of seconds that a managed Backup Exec server waits between sending job status updates to the central administration server. To preserve network bandwidth when many jobs are running, increase the amount of time between job update statuses. Decrease the amount of time if you want to send more updates.

The default is 10 seconds, which provides near real-time monitoring. This setting is recommended only for fast network connections.

Send status updates to the central administration server every

Note: This option appears only if **Custom** is selected in the **Connection settings** field and **Yes** is selected in the **Send active job status updates to the central administration server** field.

Set the amount of time that a managed Backup Exec server waits between sending job status updates to the central administration server. To preserve network bandwidth when many jobs are running, increase the amount of time between job update statuses. Decrease the amount of time if you want to send more updates.

The default is 10 seconds, which provides near real-time monitoring. This setting is recommended only for fast network connections.

For low-bandwidth network connections, consider a setting of 120 seconds. This frequency allows updates to be displayed for a medium-sized job while still significantly decreasing the network traffic caused by job status updates.

If you uncheck the check box, job status updates are not sent. Job progress is not displayed on the central administration server. When the job is complete, the **Job History** on the central administration server is updated.

Send job log details to the central administration server

Choose when the job log for the managed Backup Exec server is sent to the central administration server. You can choose to send the job log one time per day, after a job completes, or never.

The following options are available:

- **Never**
If you select this option, job logs are stored locally at the managed Backup Exec server.
- **Once a day**
If you select this option, the **Send job logs at** field appears. You must select the time to send the job log to the central administration server.
- **On job completion**
If you select this option, the **Send job log only if the job fails** field appears. Select **Yes** to send the job log only for failed jobs. Select **No** to send the job log regardless of the job disposition.

Send job logs at

Choose the time when Backup Exec sends the job logs for the managed Backup Exec server to the central administration server. This option appears only if **Once a day** is selected in the option **Send job log details to the central administration server**.

Send job log only if the job fails

Choose whether to send the job log for failed jobs only or for all jobs. Select **Yes** to send the job log only for failed jobs. Select **No** to send the job log regardless of the job disposition. This option appears only if **On job completion** is selected in the option **Send job log details to the central administration server**.

Send job history details to central administration server

Choose when the job history for the managed Backup Exec server is sent to the central administration server.

The following options are available:

- **Never**

If you select this option, job histories are stored locally at the managed Backup Exec server.

- **Once a day**

If you select this option, the **Send job history logs at** field appears. You must select the time to send the job history to the central administration server.

- **On job completion**

If you select this option, the **Send job history only if the job fails** field appears. Select **Yes** to send the job history only for failed jobs. Select **No** to send the job history regardless of the job disposition.

Send job history at

Choose the time when Backup Exec sends the job history for the managed Backup Exec server to the central administration server. This option appears only if **Once a day** is selected in the option **Send job history details to the central administration server**.

Send job history details only if the job fails

Choose whether to send job history details for failed jobs only or for all jobs. Select **Yes** to send the job history details only for failed jobs. Select **No** to send the job history details regardless of the job disposition. This option appears only if **On job completion** is selected in the option **Send job history details to the central administration server**.

Monitor jobs that are created locally on the managed Backup Exec server

Select this option if you want to view delegated jobs and the jobs that are created on the local managed Backup Exec server.

You can also hold, delete, run, cancel, and change the priority order of the jobs that are created on or copied to the local managed Backup Exec server.

Display an alert when the time is not synchronized between the managed Backup Exec server and the central administration server

Select this option to enable Backup Exec to create an alert if the clock on the managed Backup Exec server differs from the clock on the central administration server. An alert is generated when the number of seconds indicated is exceeded.

CAS monitors the internal computer clocks on both the managed Backup Exec servers and the central administration server. If time differences develop between the central administration server and the managed Backup Exec servers, jobs could run at unexpected times. To prevent problems, the time that is reported on managed Backup Exec servers should match the time that is reported on the central administration server. If you receive time difference alerts, reset the clock on the managed Backup Exec server to match the system clock on the central administration server.

If you change the system time on either the managed Backup Exec server or the central administration server, you must restart the Backup Exec services on that server.

Send the alert after the servers are not synchronized for

Indicate the number of seconds that the clocks on the managed Backup Exec server and the central administration server must differ before Backup Exec sends an alert.

Note: This option appears only if **Enabled** is selected in the **Display an alert when the time is not synchronized between the managed Backup Exec server and the central administration server** field.

Storage and media database location

Indicate whether the storage and media database is located on the central administration server or a managed Backup Exec server.

Keep the catalogs on

Set the location of the catalog to one of the following locations:

- **Managed Backup Exec server (distributed)**
Distributes the catalog files between the central administration server and the managed Backup Exec server. If storage and media data is kept in a local database on the managed Backup Exec server, then the distributed location is the only available catalog location. Select this option if you have a low-bandwidth network connection.
- **Central administration server (centralized)**
Keeps all catalog files on the central administration server. A high-bandwidth network connection is required for this option.
- **Both servers (replicated)**
Replicates all catalog files from the managed Backup Exec server to the central administration server. If a managed Backup Exec server is unavailable, you can still browse the catalog from the central administration server. However, you cannot restore data because the managed Backup Exec server is unavailable. A high-bandwidth network connection is required for this option.

See [“CAS distributed, centralized, and replicated catalog locations”](#) on page 1344.

Private cloud server

Select this option to enable a managed service provider to locate a Backup Exec server in its data center, and then configure it for a CAS environment with other Backup Exec servers that are located across the WAN at the managed service provider's customer locations. As an alternative to shipping tapes off-site for storage, backups can be run and stored locally, and then copied to the cloud server's deduplication disk storage device. Additionally, this feature can be used by customers with widely distributed networks who want to use Backup Exec servers in remote offices for local backups, and then copy the backup sets to a Backup Exec server that is located in a central data center. This option is part of the Cloud Services for Backup Exec feature.

4 Click **Apply**.

What happens when CAS communication thresholds are reached

In a CAS environment, communications that occur between managed Backup Exec servers and the central administration server can sometimes be disrupted even if network communications are normal. If job-related communication disruptions occur between a managed Backup Exec server and the central administration server, the managed Backup Exec server's communication status changes from Enabled to Stalled or No Communication. The jobs waiting to be processed by the managed Backup Exec server are held in the managed Backup Exec server's job queue until the communications are restored.

You can set the amount of time that Backup Exec waits before changing the managed Backup Exec server's status if it becomes unresponsive. When a managed Backup Exec server's status changes to Stalled or No Communication, the central administration server changes how it handles current and future jobs that are delegated to the stalled managed Backup Exec server.

For example, if communications from a managed Backup Exec server are not received at the central administration server after the set amount of time, the central administration server changes the Backup Exec server's communication

status to Stalled. Job delegation to the managed Backup Exec server is suspended as it continues to wait for the managed Backup Exec server to return to an Enabled status. Jobs are delegated to other managed Backup Exec servers that are represented in the destination storage device or Backup Exec server pool.

CAS continues to monitor the amount of time during which no communications are received from the managed Backup Exec server. After a set amount of time passes after a Stalled status appears, CAS changes the status of the managed Backup Exec server to No Communication. CAS marks the jobs as Failed, and then begins job recovery by invoking the custom error-handling rule Recovered Jobs for any job that is active at the time the No Communication status appears.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1346.

Enabling or disabling communications between the managed Backup Exec server and the central administration server

You can manually enable or disable communications between the managed Backup Exec server and the central administration server. When communications are disabled, jobs cannot be delegated to the managed Backup Exec server.

To enable communications between the managed Backup Exec server and the central administration server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to enable communications.
- 2 Select **Communication Enabled**.

To disable communications between the managed Backup Exec server and the central administration server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to disable communications.
- 2 Select **Communication Enabled** to remove the check mark.

Alerts and notifications in CAS

In a Central Admin Server feature (CAS) environment, the alerts that are generated on a managed Backup Exec server are automatically rolled up to the central administration server. To see those alerts on the central administration server, you must configure alert categories to enable or disable alerts on each managed Backup Exec server and on the central administration server itself.

After you respond to and clear the active alert on the central administration server, the alert is cleared on the managed Backup Exec server as well.

If you enable Backup Exec alerts on a managed Backup Exec server without enabling alerts on the central administration server, alerts appear only on the managed Backup Exec server where they are generated; they do not appear on the central administration server.

On the central administration server, you can view alerts for all managed Backup Exec servers, or you can filter the alerts to view only the alerts for a specific managed Backup Exec server or Backup Exec server pool.

You can configure a notification on either the central administration server or the managed Backup Exec server. Regardless of where you configure the notification, if it is for a delegated job, it is sent by the central administration server. You can choose to notify the local administrator of the managed Backup Exec server, or the administrator of the central administration server, or both.

Copying alert configurations to managed Backup Exec servers

You can enable and configure alerts at the central administration server, and then copy the alert configurations to a managed Backup Exec server. After the alert configurations are copied, the alerts that are generated on a managed Backup Exec server appear on both the managed Backup Exec server and the central administration server.

To copy alert configurations to managed Backup Exec servers

- 1 On the central administration server, click the Backup Exec button.
- 2 Select **Configuration and Settings**, and then select **Copy Settings to Other Servers**.
- 3 Under **Select settings to copy**, check **Alert configuration**.
- 4 Click **Add**.
- 5 Enter the name of a managed Backup Exec server to which you want to copy the alert configuration.
- 6 Click **OK**.
- 7 On the **Copy Settings** dialog box, click **OK**.

An alert on the central administration server confirms that the copy succeeded.

Enabling managed Backup Exec servers to use any available network interface card

By default, jobs that are delegated or copied to a managed Backup Exec server from the central administration server use the network and security settings that are set on the managed Backup Exec server.

However, you can select an option on the central administration server to let a job use any network interface to access Backup Exec agents if the selected network interface is unavailable. Enabling this option for a backup job lets the managed Backup Exec server use an alternative network interface to run important backup jobs that would otherwise fail to run.

To enable managed Backup Exec servers to use any available network interface card

- 1 On the central administration server, create a backup job.
- 2 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 3 In the left pane, select **Network**.
- 4 Check **Allow managed Backup Exec server to use any network interface to access Backup Exec agents**.
- 5 Set any additional options for the backup job
- 6 Click **OK**.

About job delegation in CAS

Job delegation is the automatic load balancing of jobs among the various storage devices that are attached to the managed Backup Exec servers. The job is created on the central administration server, but can be run on any managed Backup Exec server.

When the storage devices are logically grouped in Backup Exec server pools, and as the storage devices become available, they process jobs that are delegated from the central administration server. For example, if a storage pool contains two storage devices and one is busy processing a job, the central administration server automatically delegates another job to the idle storage device.

See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.

About copying jobs instead of delegating jobs in CAS

If the managed Backup Exec server's storage and media data are kept on a local database on the managed Backup Exec server, the central administration server cannot delegate jobs to it. Instead, you can copy job options, default schedules, error-handling rules, and alert configurations from the central administration server to the managed Backup Exec server. A persistent network connection to the central administration server is not needed when the jobs are run locally on the managed Backup Exec server.

Use the same names for objects on the central administration server and all of the managed Backup Exec servers that you want to copy jobs to. For example, use the same name for a storage pool on the central administration server and on the managed Backup Exec server. Then, it is not necessary to customize settings or names for each managed Backup Exec server that you copy jobs to.

See ["Copying configuration settings to another Backup Exec server"](#) on page 776.

About adding storage devices in a CAS environment

From the central administration server, you can run the **Configure Storage Wizard** to set up devices for the central administration server or for any of the managed Backup Exec servers. After managed Backup Exec servers are installed, they appear on the **Storage** tab of the central administration server. When you start the **Configure Storage Wizard**, you are prompted to select the server for which you want to configure storage. You can choose the central administration server or any managed Backup Exec server that runs the same version of Backup Exec as the central administration server.

How data lifecycle management (DLM) works in a CAS environment

Backup Exec uses data lifecycle management (DLM) to delete the expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, and virtual disks. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. If backup sets are dependent of other backup sets, then Backup Exec does not delete the backup set until all expiration dates on the backup sets are reached. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

If you want Backup Exec to delete all expired backup sets, even if they are the last remaining backup sets that you need to restore a server, you can select the

option **Allow Backup Exec to delete all expired backup sets** on the **Storage** settings dialog box. In a CAS environment, this option appears only on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers.

Warning: If you enable the option **Allow Backup Exec to delete all expired backup sets**, the data that you need to restore a server may not be available.

If you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 352.

Obtaining media audit information for a managed Backup Exec server

The Media Audit report lists the recent configuration changes that were made to your media. In a CAS environment, if you run this report from the central administration server the report provides data only for the media for the central administration server. It does not provide any data for any of the managed Backup Exec servers. To obtain the media audit data for a managed Backup Exec server, you must either log on to the local managed Backup Exec server or access the managed Backup Exec server from the Remote Administration Console.

To obtain media audit information for a managed Backup Exec server

- 1 Do one of the following:
 - Log on locally to the managed Backup Exec server.

- From a remote Windows server or workstation, click **Start**, point to **Backup Exec**, and then enter the name of the server you want to connect to and the credentials for that server.
- 2 On the **Reports** tab, under **Report Groups**, select **Configuration**.
- 3 Select **Media Audit** from the list of reports, and then click **Run Report Now** to run the report immediately or click **New Scheduled Report** to schedule the report to run later.

How to use Backup Exec server pools in CAS


In a CAS environment, you can group multiple managed Backup Exec servers together into Backup Exec server pools. If you create a pool of managed Backup Exec servers, all of the pools on those managed Backup Exec servers are available for job delegation. If there are multiple devices attached to each of the managed Backup Exec servers in the Backup Exec server pool, you can create multiple, smaller pools that are made up of fewer storage devices. Use this method to send some jobs to a specific pool in the Backup Exec server pool, and send other jobs to a different pool in the same Backup Exec server pool.

Backup Exec server pools can contain multiple managed Backup Exec servers or just one managed Backup Exec server. A managed Backup Exec server can belong to more than one Backup Exec server pool. The central administration server can be used as a managed Backup Exec server and can be included in the Backup Exec server pool.

Any managed Backup Exec server or Backup Exec server in a pool must be able to access the destination device for the backup. If there is no intersection between the device and the managed Backup Exec server or Backup Exec server pools, the job does not run. The Jobs list displays the following status: Ready; No Backup Exec server available in Backup Exec server pool.

This graphic shows a Backup Exec server pool.

Figure K-3 An example of a CAS-configured Backup Exec server pool inside a corporate network



See [“Creating a Backup Exec server pool”](#) on page 1364.

See [“Selecting a Backup Exec server pool for backups”](#) on page 1364.

Selecting a Backup Exec server pool for backups

You can run a job on storage devices that are on a specific managed Backup Exec server or on storage devices that are in a group of managed Backup Exec servers. This filter lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed Backup Exec servers in a pool named Exchange Backups, you can select this option, and then select the Exchange Backups server pool.

To select a Backup Exec server pool for backups

- 1 Create a backup definition.
- 2 On the **Backup** box, click **Edit**.
- 3 On the **Backup Options** dialog box, in the left pane, select **Storage**.
- 4 In **Backup Exec server or Backup Exec server pool**, select the pool that you want to use for all of the backups in the backup definition.
- 5 Configure any additional options for the backup definition.

See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.

Creating a Backup Exec server pool

You can group, or pool, Backup Exec servers.

See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.

See [“Adding managed Backup Exec servers to a Backup Exec server pool ”](#) on page 1365.

To create a Backup Exec server pool

- 1 On the central administration server's **Storage** tab, in the **Configure** group, select **Configure Storage**.
- 2 Select the server on which you want to create the server pool, and then click **Next**.
- 3 Select **Storage pools**, and then click **Next**.
- 4 Select **Backup Exec server pool**, and then click **Next**.
- 5 Enter a name and description for the pool, and then click **Next**.
- 6 Check the check boxes next to the names of the servers that you want to include in the pool, and then click **Next**.
- 7 Click **Finish**.

Adding managed Backup Exec servers to a Backup Exec server pool

You can add managed Backup Exec servers to existing Backup Exec server pools.

See [“Creating a Backup Exec server pool”](#) on page 1364.

To add managed Backup Exec servers to a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**, and then double-click a Backup Exec server pool to which you want to add managed Backup Exec servers.
- 3 Under **Backup Exec servers that belong to the pool**, click **Add**.
- 4 Check the check boxes for the Backup Exec servers that you want to add to the pool, and then click **OK**.
- 5 Click **Apply**.

Deleting a Backup Exec server pool

You can delete a Backup Exec server pool at any time.

To delete a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**.

- 3 Right-click the Backup Exec server pool that you want to delete, and then click **Delete**.
- 4 Click **Yes** to confirm that you want to delete the pool.

See [“Removing a managed Backup Exec server from a Backup Exec server pool”](#) on page 1366.

Removing a managed Backup Exec server from a Backup Exec server pool

Removing a managed Backup Exec server deletes it from a Backup Exec server pool, but does not remove it from Backup Exec.

To remove a managed Backup Exec server from a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**, and then double-click the Backup Exec server pool that contains the server you want to delete.
- 3 Under **Backup Exec servers that belong to the pool**, select the Backup Exec servers that you want to remove from the pool, and then click **Remove**.
- 4 Click **Apply**.

See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.

How centralized restore works in CAS

Depending on whether the required storage media resides in storage devices or is stored off-site, initiating restore operations from the central administration server can be an automated process with little user intervention necessary.

When you use centralized restore with online media, you run the Restore Wizard at the central administration server. During the data selection process, CAS determines which media are required to complete the restore operation, and then queries the Backup Exec storage and media database to determine the identity of the storage device where the media reside. After you run the Restore Wizard, CAS begins the restore operation by delegating the jobs to the central administration server or managed Backup Exec servers that control the selected storage devices. If the data that is being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

When you use centralized restore with offline media, you run the Restore Wizards at the central administration server. During the data selection process, CAS determines what media is required to complete the restore operation, and then

queries the Backup Exec storage and media database to determine the identity of the storage device where the primary media resides. If the media is not found in a storage device, the media is considered offline. CAS then presents you with a selection of drive pools and storage devices that are compatible with the type of media being used during the restore operation, thus giving you the flexibility of choosing a storage device in which to load your media.

After noting the identity and location of the storage device you have selected to run the job, you do the following:

- Submit the restore job on hold as a scheduled job.
- Retrieve the media, place it in the storage device.
- Remove the job from hold at the central administration server, at which time the restore job begins.

CAS then delegates the job to the managed Backup Exec server that controls the selected storage device. If the data being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

Before restore operations from the central administration server can be initiated, the following requirements must be met:

- Managed Backup Exec server communication status must be Enabled.
- Managed Backup Exec servers must be online with all Backup Exec server statuses showing Online.

Best practices for centralized restore in CAS

It is recommended that you use the following best practices for centralized restore:

- Select only one resource to restore for each job.
- Select the same restore device or Backup Exec server for all of the selections that are in the same restore job.
- Select a Backup Exec server that has compatible devices for all media that is required for the restore job.

How CAS restores data that resides on multiple storage devices

If the data selected for restore is located on a single device that is attached to a managed Backup Exec server, then a single restore job is created at, and then delegated from, the central administration server. However, if the data being selected for restore is located on multiple devices in the CAS environment, then the single restore job is split into separate restore jobs, depending on the number of devices involved.


All split restore jobs have the same name as the original job, but are differentiated and linked with a subscript numeral that is appended to the job name.

For example, if you create a restore job and the data you select for restore resides in one device on a managed Backup Exec server, CAS creates one restore job. However, if you create one restore job and the data you select resides on two or more devices that are attached to a managed Backup Exec server, CAS creates two or more restore jobs.


The following graphic shows how CAS restores data that is stored on a single device.

Figure K-4 For data stored on a single storage device

You select the data to restore from the administration console of the central administration server.




CASO creates a restore job, which is delegated for processing at the managed Backup Exec server where the data resides.




The following graphic shows how CAS restores data that is stored on multiple devices.

Figure K-5 For data stored on multiple storage devices

You select the data to restore from the administration console of the central administration server.



CASO creates multiple restore jobs, which are delegated for processing at the managed Backup Exec servers where the data resides.



About recovering failed jobs in CAS

The Backup Exec error-handling rule named Recovered Jobs is a custom error-handling rule that is used by CAS to recover the jobs that failed because of issues with internal job communications. This rule is created when Backup Exec is installed and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. During the first retry attempt, CAS attempts to re-delegate the jobs to another available managed Backup Exec server.

If this attempt fails, CAS makes a second attempt at finding another available managed Backup Exec server to process the jobs. If no managed Backup Exec servers are available, the final job disposition is to place the job on hold until you have fixed the error condition.

Note: If the checkpoint restart error-handling rule is enabled, then recovered jobs are not resubmitted to a Backup Exec server pool to be run on a different server. The checkpoint restart error-handling rule reschedules the job to run on the original server when that server comes online. To enable a recovered job to be resubmitted to a Backup Exec server pool, you must disable the checkpoint restart error-handling rule.

Note: If you target a job to a Backup Exec server pool that contains multiple managed Backup Exec servers and a job failure occurs, the recovery process uses only the managed Backup Exec servers in the Backup Exec server pool. Managed Backup Exec servers that are not in the Backup Exec server pool are not used for job recovery.

When you open the job history entry for a recovered job, the reason for the failure is listed as Job Errors, with an explanation of the type of internal communication error that occurred. The job history entry also indicates that the job was recovered.

Note: Job logs are not created for the jobs that are recovered.

The following table describes the CAS error codes that are selected by default for the Recovered Jobs custom error-handling rule:

Table K-3 Error codes for Recovered Jobs custom error-handling rule

Error code	Description
0xE000881B JOBDISPATCH	The displayed message is: Job failed while being dispatched. The job will be recovered.
0xE000881D JOB_CASO_QUEUE FAILURE	The displayed message is: The job could not be delegated to the destination managed Backup Exec server. The managed Backup Exec server may not be online, or there may be a communications failure. The job will be recovered.
0xE000881E JOB_CASO_REMOTEMMS_STARTFAILURE	The displayed message is: The job failed to start on the destination managed Backup Exec server, possibly because a database error occurred. The job will be recovered.

See [“Error-handling rules for failed or canceled jobs”](#) on page 285.

See [“Custom error-handling rule for recovered jobs”](#) on page 288.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1346.

Pausing or resuming a managed Backup Exec server

You can pause and resume a managed Backup Exec server from the central administration server.

Pausing a managed Backup Exec server prevents the central administration server from delegating jobs to it. When the managed Backup Exec server is paused, its status changes from Online to Paused.

Caution: When you install Backup Exec features at a managed Backup Exec server, the managed Backup Exec server must be paused so that no further jobs are delegated to it from the central administration server while the installation process occurs. If jobs are running, let them finish or cancel them before beginning the installation.

When you resume a paused managed Backup Exec server, the following changes occur:

- Jobs can be delegated from the central administration server to the managed Backup Exec server.
- The managed Backup Exec server's status changes from Paused to Online in the **State** column on the **Storage** tab.

To pause a managed Backup Exec server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to pause.
- 2 Click **Pause**.

To resume a paused managed Backup Exec server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to resume.
- 2 Click **Pause** to remove the check mark next to **Pause**.

Stopping or starting Backup Exec services for a managed Backup Exec server

You can stop or start the Backup Exec services on a managed Backup Exec server from the central administration server.

To stop Backup Exec services for a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to stop services.
- 2 Select **Backup Exec Services**.
- 3 On the **Backup Exec Services Manager** dialog box, click **Stop all services**.
- 4 Click **Close**.

To start Backup Exec services for a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to start services.
- 2 Select **Backup Exec Services**.
- 3 On the **Backup Exec Services Manager** dialog box, click **Start all services**.
- 4 Click **Close**.

Viewing managed Backup Exec server properties

From the central administration server, you can view properties for managed Backup Exec servers.

To view managed Backup Exec server properties

- 1 On the central administration server's **Storage** tab, double-click the managed Backup Exec server for which you want to view properties.

2 In the left pane, select **Properties**.

Name	Displays the name of the managed Backup Exec server or the central administration server.
Description	Indicates whether the server is a managed Backup Exec server or a central administration server. You can change this description.
Server status	Indicates the current status of the server, such as online, paused, unavailable, or offline.
Version	Indicates the version of Backup Exec that is installed.
License	Provides information about the Backup Exec license that is installed on the server.
License usage compliant	Indicates if the Backup Exec server is compliant with the license installed.
Time zone	Indicates the time zone that is set for the server.
Start date and time	Indicates when the server was started.
Current date and time	Indicates the current date and time on the server.
Operating system type	Indicates the type of operating system that is installed on the server.
Operating system version	Indicates the version of the operating system that is installed on the server.
Operating system build	Indicates the build number of the operating system that is installed on the server.
Processor type	Indicates the type of processor that the server has.
Number of processors	Indicates the number of processors that the server has.
Total physical memory	Indicates the total amount of physical memory that the server has.

Available physical memory	Indicates the amount of physical memory that is available on the server.
Total virtual memory	Indicates the total amount of virtual memory that the server has.
Available virtual memory	Indicates the amount of virtual memory that is available on the server.
Total pagefile size	Indicate the total amount of memory that is available in the server's page file.

Viewing the settings for a central administration server

If you have the Central Admin Server feature (CAS), you can view information about the location of the databases for Backup Exec. The databases include the Backup Exec Database, the Advanced Device and Media Database (ADAMM), and the catalog database.

During Backup Exec installation, if you chose the default option to create a local Backup Exec SQL Express instance on which to store the Backup Exec Database, the databases are all located on the local Backup Exec server. If you chose another instance on the network on which to store the Backup Exec Database, then the databases are all located on the Microsoft SQL Server that contains that instance.

To view the settings for a central administration server

- 1 Do one of the following:
 - Click the Backup Exec button, select **Configuration and Settings**, and then click **Local server properties**.

- On the **Storage** tab, double-click the central administration server.
- 2 In the left pane, click **Settings**.

Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec Database.
Instance	Shows the name of the instance that the Backup Exec Database is installed on.
Name	Shows the Microsoft SQL Server name of the Backup Exec Database.
Path	Shows the path of the Backup Exec Database.
Server	Shows the name of the Microsoft SQL Server that contains the ADAMM database.
Instance	Shows the name of the instance that the ADAMM database is installed on.
Name	Shows the Microsoft SQL Server name for the ADAMM database.
Path	Shows the path of the ADAMM database.
Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec catalog database.
Instance	Shows the database instance that contains the catalog database.
Name	Shows the Microsoft SQL Server name for the Backup Exec catalog database.
Path	Shows the path of the Backup Exec catalog database.
Private cloud server	Indicates if the private cloud server option is enabled or disabled.

Disaster recovery in CAS

Use the Backup Exec Simplified Disaster Recovery (SDR) feature to protect both managed Backup Exec servers and the central administration server in a CAS environment.

See [“About Simplified Disaster Recovery”](#) on page 877.

Before implementing the SDR feature in a CAS environment, review the following:

- To create recovery media for any managed Backup Exec server or central administration server, the **Create Simplified Disaster Recovery Disk Wizard** must be run at the central administration server. If you use a remote administration environment, connect to the central administration server.
- If you want managed Backup Exec servers to be protected using a bootable disk image, you must run the **Create Simplified Disaster Recovery Disk Wizard** at each of the managed Backup Exec servers where a bootable disk device is installed.
- You must locally back up and restore a central administration server.

Disaster recovery of a managed Backup Exec server that is managed by a failed central administration server

If communications are still active between the managed Backup Exec server and the central administration server, you can change a managed Backup Exec server to a standalone Backup Exec server by deleting it from the **Storage** tab. However, if the central administration server experiences a disaster, communications are lost between the managed Backup Exec server and the central administration server and the managed Backup Exec server cannot be deleted from the **Storage** tab. However, you can change the managed Backup Exec server to a standalone Backup Exec server by using the Windows Change Program feature.

To change a managed Backup Exec server to a standalone Backup Exec server when communications between the two servers have been lost

- 1 On the managed Backup Exec server, open the Windows control panel.
- 2 Select **Add/Remove Programs** or **Programs and Features**, depending on the version of Windows you are using.
- 3 Select **Backup Exec** from the list, and then click **Change**.
- 4 On the **Additional Options** panel, click **Next**.
- 5 On the **Add Licenses** panel, click **Next**.
- 6 On the **Configure Options** panel, click **Next**.
- 7 On the **Choose Languages** panel, click **Next**.

- 8 On the **Central Admin Server** panel, select **Locally managed Backup Exec server**, and then click **Next**.
- 9 When the error message about the managed Backup Exec server being unable to contact the central administration server appears, click **OK**.

Note: If an error message does not appear, then Backup Exec detected a server with the same name as the central administration server on the network. If that server is the central administration server and it is running correctly, you should be able to change it to a managed Backup Exec server from the **Storage** tab on the central administration server.

- 10 Click **Next**.
- 11 On the **Installation Review** panel, click **Install**.
- 12 When the installation is complete, on the Backup Exec server, add any storage devices that may be needed.

Note: You may also need to inventory and catalog the storage device, depending on which catalog method was in use when the managed Backup Exec server was joined to the central administration server.

When the central administration server is online again, you can change this standalone Backup Exec server back to a managed Backup Exec server.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1338.

Troubleshooting CAS

If you encounter issues with CAS, review the following questions and answers.

Table K-4 Troubleshooting CAS

Question	Answer
I received error 1065 that says "Database specified does not exist". What causes this error?	This error can occur for the following reasons: ■ UDP traffic is blocked on the network between the central administration server and the managed Backup Exec server. ■ The SQL configuration on the central administration server is not set correctly. ■ When the central administration server is installed to a named SQL instance and the SQL browser service is not running. ■ The Named Pipes or TCP/IP protocols are not enabled or are not set for remote connections.
I changed the system time, but the change hasn't gone into effect on my managed Backup Exec servers or central administration server. Why?	If you change the system time on either the managed Backup Exec server or the central administration server, you must restart the Backup Exec services. Backup Exec processes the time change when the services restart.
I received the error "Ready, Job storage does not contain any eligible devices". What causes this error?	This error can occur if you remove managed Backup Exec servers from the wrong location. To remove managed Backup Exec servers from a central administration server, use the Delete option on the Storage tab. When you remove a managed Backup Exec server from the Storage tab, the server becomes a standalone server, so it can still be backed up and restored as a standalone server, but it can no longer be controlled by the central administration server. Do not remove the managed Backup Exec server from the Backup and Restore tab. See "Deleting a managed Backup Exec server from a CAS environment" on page 1339.

See ["About the Central Admin Server feature"](#) on page 1321.

Running the Backup Exec Utility for CAS operations

A separate application called Backup Exec Utility is available to help you perform the following CAS operations:

- Move a managed Backup Exec server.
- Disable or enable communication with a managed Backup Exec server.

Use the Backup Exec Utility only with the guidance of Technical Support. Improper use of this utility can result in configuration changes that may prevent Backup Exec from running.

To run the Backup Exec Utility

- 1 From the Backup Exec installation directory, in `<Backup Exec install path>\Backup Exec`, double-click **BEUtility**.
- 2 On the **Backup Exec Utility** menu, click **Help** to learn about how to use BE Utility.

Uninstalling Backup Exec from the central administration server

Before you uninstall Backup Exec from the central administration server, you must delete all managed Backup Exec servers from the **Storage** tab on the central administration server.

Caution: Failure to uninstall in the following sequence may result in long delays when shutting down Backup Exec services during the uninstall of Backup Exec on the managed Backup Exec servers.

To uninstall Backup Exec from the central administration server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server.
- 2 Select **Delete**.
- 3 Click **Yes** to confirm that you want to delete.
- 4 Repeat steps 1 through 3 for each managed Backup Exec server that the central administration server manages.
- 5 Uninstall Backup Exec from the central administration server.

See ["Uninstalling Backup Exec"](#) on page 110.

Uninstalling Backup Exec from a managed Backup Exec server

You must delete the managed Backup Exec server from the **Storage** tab on the central administration server before you uninstall Backup Exec.

To uninstall Backup Exec from a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server.
- 2 Select **Delete**.
- 3 Click **Yes** to confirm that you want to delete.
- 4 Uninstall Backup Exec from the managed Backup Exec server.

See ["Uninstalling Backup Exec"](#) on page 110.

Backup Exec Advanced Disk-based Backup Feature

This appendix includes the following topics:

- [About the Advanced Disk-based Backup feature](#)
- [How to use synthetic backups in place of recurring full backups](#)
- [Setting default backup options for the Advanced Disk-based Backup feature](#)
- [About true image restore for synthetic backups](#)
- [How to use off-host backup to process remote computer backups on the Backup Exec server](#)
- [Configuring off-host backup options for a backup job](#)
- [Best practices for off-host backup](#)
- [Troubleshooting off-host backups](#)
- [Off-host backup issues with hardware providers](#)

About the Advanced Disk-based Backup feature

The Advanced Disk-based Backup feature provides the following features:

- **Synthetic backup**
This feature enables a full backup to be assembled, or synthesized, from a baseline full backup and subsequent incremental backups.

The benefits of using a synthetic backup include the following:

- A reduced backup window, since the synthetic backup can be scheduled outside of the time-critical backup window.
- Reduced network traffic, since the synthetic backup does not need to access the network.
- True image restore, which enables Backup Exec to restore the contents of directories to what they were at the time of any synthetic full or incremental backup.
- Off-host backup
This feature enables the backup operation to be processed on a Backup Exec server instead of on the remote computer or host computer. Moving the backup from the remote computer to the Backup Exec server enables better backup performance and frees the remote computer as well.

The Advanced Disk-based Backup feature (ADBO) is installed as part of the Enterprise Server feature, for which you must enter a license on the Backup Exec server.

See [“How to use synthetic backups in place of recurring full backups”](#) on page 1383.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1388.

See [“About the Backup Exec installation process and licensing options”](#) on page 39.

See [“Setting default backup options for the Advanced Disk-based Backup feature”](#) on page 1385.

How to use synthetic backups in place of recurring full backups

The synthetic backup feature eliminates the need to perform recurring full backups for supported remote resources. The synthetic backup is assembled from a full backup (called a baseline) and subsequent incremental backups.

The resulting synthetic backup then becomes the new baseline. Only incremental backups are required until the next synthetic backup is created. The synthetic backup is as current as the last incremental backup that it contains.

The components of a synthetic backup are as follows:

- Baseline backup.
The baseline backup is the first full backup to run that is associated with the synthetic backup. The full baseline backup runs one time only, and backs up all of the files on the selected computer when it runs.

- **Recurring incremental backups.**
Incremental backup jobs back up the files that change after the baseline backup.
- **Recurring synthetic backups.**
The synthetic backup process combines the data from the baseline backup and the incremental backups to form a synthesized full backup of the selected computer. This synthesized full backup becomes a new baseline backup, which is combined with subsequent incremental backup sets to form a new synthesized full backup.

For any of the backups in a synthetic backup, you can add a stage to duplicate the backup data to tape.

True image restore is automatically enabled for synthetic backups. True image restore lets you restore directories as they existed at the time of the synthetic backup. Files that were deleted before the time of the synthetic backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate synthetic full or incremental backups that contain them.

Only file system data is supported for synthetic backup. Supported data includes common file system objects, such as volumes, drives, and folders.

Requirements for synthetic backup

Before you create a synthetic backup, review the following information:

- If you use an encryption key, all of the associated backups must use the same encryption key. Do not change the encryption key after the backups are created. The encryption key that is selected in the associated backups is automatically applied to the synthetic backup.
- You must configure disk storage before you can create a synthetic backup. For synthetic backups, incremental backups must use disk storage. The baseline full backup and the synthetic full backup can use tape or disk storage. See [“Configuring disk storage”](#) on page 334.
- If you send the baseline backup job to tape storage, and you want to use tape storage for the synthetic backup job, you must have two tape drives. You must use a tape drive to mount the baseline backup, and use a tape drive to mount the synthetic backup.
- You can select only file system data for synthetic backup.

See [“Configuring off-host backup options for a backup job”](#) on page 1391.

See [“About true image restore for synthetic backups”](#) on page 1386.

Setting default backup options for the Advanced Disk-based Backup feature

You can use the defaults that Backup Exec sets during installation for all Advanced Disk-based Backup feature backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Advanced Disk-based Backup feature

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for Advanced Disk-based Backup feature backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select **Advanced Disk-based Backup**.
- 4 Select the appropriate options:

Use offhost backup to move backup processing from the remote computer to the Backup Exec server

Select this option to enable off-host backup.

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

Continue the backup job (offhost backup is not used)

Select this option to let the backup job complete without using the off-host feature if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Fail the backup job (further selections are not backed up after failure occurs)

Select this option to fail the backup job if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Process logical volumes for offhost backup one at a time

Select this option to enable the backup of multiple volumes in one job, while a snapshot is created of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time.

After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

See [“About the Advanced Disk-based Backup feature”](#) on page 1382.

See [“Backing up data”](#) on page 158.

About true image restore for synthetic backups

True image restore is automatically enabled for synthetic backups. True image restore enables Backup Exec to restore the contents of directories to what they were at the time of any full backup or incremental backup. Restore selections in backup sets are made from a view of the directories as they existed at the time of the synthetic backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not restored and then overwritten.

Note: The **Use storage-based catalogs** option must be enabled in the **Catalogs** settings to use true image restore.

See [“Configuring default options for catalogs”](#) on page 255.

Backup Exec collects the information that is required to detect the files and directories that have been moved, renamed, or newly installed from a tape archive (tar) or a compressed archive. Depending on how the files were packaged and how they were installed, some newly installed files are not backed up by normal incremental backups. With true image restore enabled, Backup Exec compares path names with path names from the previous full or incremental backup. If a name is new or changed, the file or directory is backed up.

The following are examples where using true image restore backs up the files that would not otherwise be backed up:

- A file named C:\pub\doc is moved to or installed in C:\spec\doc. Here, the archive bit is unchanged for files and subdirectories inside that directory, but C:\pub\doc is new in the C:\spec\ directory and is backed up.
- A directory named C:\security\dev\ is renamed as C:\security\devices\. Here, the archive bit is unchanged for files and subdirectories inside that directory, but C:\security\devices\ is a new directory and is backed up.

The following table lists the files that are backed up in the C:\user\doc directory during a series of backups between December 1, 2012 and December 4, 2012:

Table L-1

Example table of files backed up because true image restore is enabled

Day	Type of backup	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc
December 1, 2012	Full	file1	file2	dirA\fileA	dirB\fileB	file3	
December 2, 2012	Incremental	file1	file2	dirA\fileA	-----	-----	
December 3, 2012	Incremental	file1	file2	dirA\fileA	-----	-----	
December 4, 2012	Incremental	file1	file2	-----	-----	-----	file4

Note: Dashes (-----) indicate that the file was deleted before this backup.

Assume that you want to restore the December 4, 2012 version of the C:\user\doc directory.

You perform a regular restore of the full backup set followed by a regular restore of subsequent incremental backup sets. The restored directory contains all files and directories that ever existed in C:\user\doc from December 1, 2012 (last full backup) through December 4, 2012.

For example, the following files and directories are included:

- file1
- file2
- dirA\fileA
- dirB\fileB

- file3
- file4

In a true image restore of the December 4, 2012 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on December 4, 2012.

The following list includes the files and directories that existed:

- file1
- file2
- file4

Backup Exec does not restore any of the files that were deleted before the December 4, 2012 incremental backup.

The restored directory does not include the 'dirA' subdirectories, even though they were backed up on December 4, 2012. Backup Exec does not restore these directories because they did not exist at the time of the incremental backup, which was the reference for the true image restore.

A true image restore preserves the files that are currently in the directory but were not present when the backup was completed. Assume that you created a file named file5 after the incremental backup that occurred on December 4, 2012, but before doing the restore.

In this case, the directory contains the following files after the restore:

- file1
- file2
- file4
- file5

See ["How to use synthetic backups in place of recurring full backups"](#) on page 1383.

See ["Configuring off-host backup options for a backup job"](#) on page 1391.

How to use off-host backup to process remote computer backups on the Backup Exec server

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

After the backup, the snapshots are exported from the Backup Exec server and mounted back on the remote computer and resynchronized with the source volume. This process requires solutions from the hardware providers that can support transportable snapshots. Transportable snapshots are snapshots you can import to and export from the Backup Exec server. The Microsoft Volume Shadow Copy Services (VSS) provider that you select is used for each volume in the off-host backup. An off-host backup job is performed on one remote computer at a time.

Off-host backup supports the following:

- Microsoft Volume Shadow Copy Service (VSS).
- Backups for the NTFS volumes that use the full, incremental, and differential backup methods.
- SQL Agent backups for Microsoft SQL Server 2000 and later databases.
- Exchange Agent backups for Microsoft Exchange Server instances: Support for the option to use Backup Exec Granular Recovery Technology for Exchange Agent backups is included.

Requirements for off-host backup

Additionally, the following are requirements for off-host backup:

Table L-2 Off-host backup requirements

Item	Description
Backup Exec server	The Advanced Disk-based Backup feature must be installed.
Remote computer	The Agent for Windows must be installed on the remote computer.
Backup Exec server and the remote computer	The following must be installed on both the Backup Exec server and on the remote computer: ■ The same operating system. ■ The most recent Volume Shadow Copy Services (VSS) patches. ■ The Microsoft VSS hardware snapshot provider that you want to use. Otherwise, the snapshots of the volumes cannot be exported to the Backup Exec server. ■ Ability to access the SAN shared storage or logical disk that is shared between the Backup Exec server and the remote computer.

Table L-2 Off-host backup requirements (*continued*)

Item	Description
GRT-enabled off-host backup of Exchange Server resources	Off-host backup supports Exchange Agent backups for Microsoft Exchange Server. Both the Backup Exec server and the Exchange server must be running the same version with the same software patch level (including VSS patches). The devices must also be listed on the Hardware Compatibility List.
Central Admin Server feature	If the Central Admin Server feature (CAS) is installed, do not let the central administration server delegate the job. It can delegate the job to a managed Backup Exec server that does not have the off-host capability. You must manually select the storage device for the CAS jobs that use the off-host backup method.

Advanced Disk-based feature off-host backup does not support the following:

- The option **Checkpoint Restart**.
See [“Configuring checkpoint restart”](#) on page 672.
- Volumes that run Windows BitLocker Drive Encryption.
- The backup method for files **Using catalogs**.
See [“How Backup Exec determines if a file has been backed up”](#) on page 198.

Best practices for off-host backup

The following best practices are recommended:

- Do not allow source volumes and snapped volumes to share the same physical disks. Otherwise, any attempt to split the snapshot volume from the original volume fails.
- Most hardware and software providers have some limitation about the types of volumes that are transportable. Therefore, it is recommended that you use off-host backup jobs only for backing up data for which all dependent volumes can be imported and exported.
- Ensure that the volume that you select to back up can be imported and exported and that the VSS hardware provider is on the compatibility list. Otherwise, the off-host backup fails. You can choose to continue the backup if the off-host backup fails.
You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.
- The Hitachi Raid Manager log cannot be on a volume that is being snapped. Hitachi executes I/O to its Raid Manager log file during the snapshot commit

process, and the VSS coordinator blocks I/O to any drive being snapped. Therefore, if the log directory for Raid Manager is on the volume that is being snapped, then log I/O is blocked and the snap process is deadlocked.

- If the Central Admin Server feature (CAS) is installed, you must manually select the storage for the off-host backup. Otherwise, the job may be delegated to a Backup Exec server that does not have off-host capability. See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.
- When you run an off-host backup that uses a VSS hardware provider in a Microsoft Cluster (MSCS) environment, the Backup Exec server and the remote computer must not be in the same cluster group. The cluster applications cannot support the devices’ logical unit numbers (LUNs) that have duplicate signatures and partition layouts. The snapshots containing the LUNs must be transported to a host computer that is outside the cluster.

See [“Backing up data”](#) on page 158.

See [“Configuring off-host backup options for a backup job”](#) on page 1391.

See [“Setting default backup options for the Advanced Disk-based Backup feature”](#) on page 1385.

See [“Troubleshooting off-host backups”](#) on page 1394.

Configuring off-host backup options for a backup job

Off-host backup enables Backup Exec to move backup processing from the host computer, which is the remote computer that contains the volumes selected for backup, to the Backup Exec media server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the media server, where they are backed up.

To configure off-host backup options for a backup job

- 1 On the **Backup and Restore** tab, right-click the server, servers, or server group that you want to back up.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.
- 5 Click **OK**.

- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, select **Advanced Disk-based Backup**.
- 9 Set any of the following options for this job:

Use offhost backup to move backup processing from the remote computer to the Backup Exec server

Select this option to enable off-host backup.

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

Continue the backup job (offhost backup is not used)

Select this option to let the backup job complete without using the off-host feature if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Fail the backup job (further selections are not backed up after failure occurs)

Select this option to fail the backup job if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Process logical volumes for offhost backup one at a time

Select this option to enable the backup of multiple volumes in one job, while a snapshot is created of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time.

After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

- 10 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
 - 11 Click **OK**.
 - 12 On the **Backup Definition Properties** dialog box, click **OK**.
- See [“About the Advanced Disk-based Backup feature”](#) on page 1382.
- See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1388.

Best practices for off-host backup

The following best practices are recommended:

- Keep source volumes and snapped volumes from sharing the same physical disks. Otherwise, any attempt to split the snapshot volume from the original volume fails.
- Most hardware and software providers have some limitation about the types of volumes that are transportable. Therefore, it is recommended that you use off-host backup jobs only for backing up data for which all dependent volumes can be imported and exported.
- The off-host backup fails if any one volume that you select for backup cannot be imported or exported. The off-host backup also fails if the required VSS hardware provider is not on the Hardware Compatibility List. You can choose to continue the backup if the off-host backup fails.
You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.
- The Hitachi Raid Manager log cannot be on a volume that is being snapped. Hitachi executes I/O to its Raid Manager log file during the snapshot commit process, and the VSS coordinator blocks I/O to any drive being snapped. Therefore, if the log directory for Raid Manager is on the volume that is being snapped, then log I/O is blocked and the snap process is deadlocked.
- If the Central Admin Server feature (CAS) is installed, you must manually select the storage for the off-host backup. Otherwise, the job may be delegated to a Backup Exec server that does not have off-host capability.
See [“How to use Backup Exec server pools in CAS ”](#) on page 1363.
- When you run an off-host backup that uses a VSS hardware provider in a Microsoft Cluster (MSCS) environment, the Backup Exec server and the remote computer must not be in the same cluster group. The cluster applications cannot support the devices' logical unit numbers (LUNs) that have duplicate

signatures and partition layouts. The snapshots containing the LUNs must be transported to a host computer that is outside the cluster.

See [“Troubleshooting off-host backups”](#) on page 1394.

Troubleshooting off-host backups

Off-host backup requires that the VSS providers and the volumes that are to be transported are set up correctly. Not all arrays are supported with the Advanced Disk-based feature.

You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.

To troubleshoot off-host backup issues, it is recommended that you use the tools that are available from the VSS provider to verify the required setup for off-host backup.

The minimum setup requirements are as follows:

- Volumes that you want to back up are snappable.
- Volumes are shared between the remote computer and the Backup Exec server.
- An off-host backup job can only contain the volumes that can be transported to the Backup Exec server for backup.

Other factors to consider are as follows:

- A Windows operating system must be installed on both the Backup Exec server and on the remote computer. Both computers must have the most recent Volume Shadow Copy Services (VSS) patches.
- Microsoft XML Core Services (MSXML 6.0 SP2) must be installed and running on both the Backup Exec server and the remote computer.

Troubleshooting off-host backup issues depends to an extent on the VSS provider that is used for the snapshots.

The following setup issues that are common to all providers may cause off-host backup to fail:

Table L-3 Common setup issues for off-host backup

Issue	Solution
The volumes are not shared.	You must ensure that all of the volumes reside on the disks that are shared between the remote computer and the Backup Exec server. If the volumes are not shared, the import operation fails. You may need to clean up the snapshots and resynchronize the volumes manually.
The VSS provider is not installed on the Backup Exec server and the remote computer.	The provider that is used for the snapshot must be installed on both the Backup Exec server and on the remote computer. If the provider is not installed on the Backup Exec server, the import operation fails. You may need to clean up the snapshots and resynchronize the volumes manually.
All volumes are not transportable.	All of the volumes that are selected for backup must be transportable to the Backup Exec server. If Microsoft SQL, Microsoft Exchange, or other database applications are selected for backup, make sure that the databases and log files reside on transportable volumes.
The VSS provider cannot snap all of the selected volumes.	All of the volumes that are selected for backup must be transportable to the Backup Exec server. All volumes that you select for backup must be snappable by the same provider. You must ensure that the same VSS provider supports all of the volumes in a backup job.
The log path location is incorrect.	If the provider or the supporting application creates log files during normal snapshot operation, the log files should not reside on any of the volumes that are being snapped. VSS cannot flush the write buffers, and the snapshot times out. Change the log path to another volume.
The provider or VSS services are not started.	The provider service should be running and the Microsoft Windows "Volume Shadow Copy" service should not be disabled.
The credentials are incorrect.	The machine-level credentials that are used for the job should be the same on both the Backup Exec server and the remote computer. Incorrect credentials cause snapshots or the backup to fail.

Table L-3 Common setup issues for off-host backup *(continued)*

Issue	Solution
The VSS provider is not installed on all Backup Exec servers in a Central Admin Server feature (CAS) environment.	If you configure a backup job in a CAS environment, you must send the job to managed Backup Exec servers on which the selected VSS provider is installed. You should not let the central administration server delegate the job. Otherwise, the job may be delegated to a managed Backup Exec server that does not have off-host capability.
The Backup Exec server and the remote computer are in the same cluster group.	For an off-host backup in a Microsoft Cluster environment, the Backup Exec server and the remote computer must not be in the same cluster group. The cluster applications cannot support the devices' logical unit numbers (LUNs) that have duplicate signatures and partition layouts. Therefore, you must transport the snapshots that contain the LUNs to a Backup Exec server that is outside the cluster in which the host cluster resides. If you use a Hitachi 9970 and attempt to back up Microsoft Cluster data by using the Advanced Disk-based Backup feature, you may receive the following error message: The job failed with the following error: A failure occurred querying the Writer status. To correct this problem, ensure that the RM Shadow Copy Provider for Volume Snapshot Service is present and running. If the service is not running, run RMVSSPRV.exe from c:\horcm\tool. If the service is still not running, contact Hitachi for support.

Off-host backup issues with hardware providers

Hardware disk array vendors may support VSS snapshots and the transporting of volumes to the Backup Exec server for backup in a SAN environment. Using hardware providers requires a sound understanding of how disk arrays are configured for shared access between the remote computer and the Backup Exec server in a SAN.

Consult the documentation for your hardware disk array on how to set up such disk arrays for off-host backup . Specifically, note any limitations on using the disk arrays in context with VSS snapshots, and note how to verify that the volumes are transportable. It is recommended that you make use of any tools that vendors provide to help verify the setup and to troubleshoot issues.

An off-host backup issue that can occur when you use Hitachi hardware may be that Hitachi supports only basic disks for off-host backup. The off-host backup

feature is not supported if a computer uses a combination of dynamic and basic disks, and the Hitachi provider.

See [“About the Advanced Disk-based Backup feature”](#) on page 1382.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1388.

Off-host backup issues with hardware providers

Hardware disk array vendors may support VSS snapshots and the transporting of volumes to the Backup Exec server for backup in a SAN environment. Using hardware providers requires a sound understanding of how disk arrays are configured for shared access between the remote computer and the Backup Exec server in a SAN.

Consult the documentation for your hardware disk array on how to set up such disk arrays for off-host backup. Specifically, note any limitations on using the disk arrays in context with VSS snapshots, and note how to verify that the volumes are transportable. It is recommended that you make use of any tools that vendors provide to help verify the setup and to troubleshoot issues.

An off-host backup issue that can occur when you use Hitachi hardware may be that Hitachi supports only basic disks for off-host backup. The off-host backup feature is not supported if a computer uses a combination of dynamic and basic disks, and the Hitachi provider.

See [“Troubleshooting off-host backups”](#) on page 1394.

Backup Exec NDMP Feature

This appendix includes the following topics:

- [Features of the NDMP feature](#)
- [Requirements for using the NDMP feature](#)
- [About installing the NDMP feature](#)
- [Adding NDMP servers to Backup Exec](#)
- [Sharing the tape drives on NDMP servers between multiple Backup Exec servers](#)
- [NDMP backup options for NDMP servers](#)
- [NDMP server backup selections](#)
- [How to use patterns to exclude files and directories from an NDMP server backup](#)
- [Supported configurations for duplicating data from NDMP servers](#)
- [About restoring and redirecting restore data for NDMP servers](#)
- [Setting the default backup options for the NDMP feature](#)
- [Viewing the properties of an NDMP server](#)
- [Viewing storage properties for an NDMP server](#)
- [About support for NetApp's Data ONTAP](#)

Features of the NDMP feature

The Backup Exec NDMP feature uses the Network Data Management Protocol (NDMP) to back up and restore network-attached storage NDMP servers.

You can use the following configurations to back up data from a NDMP server:

- A direct-attached configuration in which a tape device or a virtual tape library is directly connected to the NDMP server.
- A three-way configuration in which a tape device or a virtual tape library is connected to another NDMP server.
- A remote configuration in which a storage device is attached to the Backup Exec server.
- A configuration in which any storage device is attached to the same SAN as the Backup Exec server.

Note: You cannot back up data from an NDMP server to a simulated tape library or to a tape device that is attached to a Backup Exec Remote Media Agent for Linux.

You cannot redirect backup sets from an NDMP server to a computer that runs the Windows or Linux operating systems.

For information about the best practices to use Backup Exec NDMP Option, refer to *Backup Exec Best Practices*.

See [“Requirements for using the NDMP feature”](#) on page 1399.

See [“Sharing the tape drives on NDMP servers between multiple Backup Exec servers”](#) on page 1402.

Requirements for using the NDMP feature

To use the NDMP feature, the following items are required:

- Backup Exec must be installed on a Windows Server that runs on a supported platform.
You can find a list of supported operating systems in the Backup Exec Software Compatibility List.
- The network-attached storage NDMP server must run version 4 of the Network Data Management Protocol.

You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List.

See [“About installing the NDMP feature”](#) on page 1400.

See [“Adding NDMP servers to Backup Exec”](#) on page 1400.

About installing the NDMP feature

The NDMP feature is installed locally on the Backup Exec server as a separate add-on component of Backup Exec. No files are copied to the network-attached storage NDMP server.

See [“Installing additional agents and features to the local Backup Exec server”](#) on page 58.

See [“Requirements for using the NDMP feature”](#) on page 1399.

Adding NDMP servers to Backup Exec

You can add a network-attached storage NDMP server to Backup Exec to back it up and to use the storage devices that are attached to it.

If you add an NDMP server in a Central Admin Server feature environment, you must add it to one of the following:

- The central administration server.
- A managed Backup Exec server on which the storage and media database is located.

Table M-1 Methods to add an NDMP server to Backup Exec

Task	Method
To add an NDMP server for backup	Use the Add Server wizard and add the NDMP server as a file server or NDMP data server. The NDMP server is added to the list of servers on the Backup and Restore tab. See “About the list of servers on the Backup and Restore tab” on page 150.

Table M-1 Methods to add an NDMP server to Backup Exec *(continued)*

Task	Method
To add an NDMP server for backup and use its attached storage devices	Use the Configure Storage wizard to add the NDMP server to the list of servers on the Backup and Restore tab. Then, after the Backup Exec services are restarted, Backup Exec automatically discovers any storage devices that are attached to the NDMP server. The storage devices appear on the Storage tab. See “To add the NDMP server for backup and use its attached storage devices” on page 1401.

To add the NDMP server for backup and use its attached storage devices

- On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- Do one of the following:

If the Central Admin Server feature is not installed

Click **Network storage**, and then click **Next**.

If the Central Admin Server feature is installed

Do the following in the order listed:
 - Select the Backup Exec server on which you want to configure storage, and then click **Next**.
 - Click **Network storage**, and then click **Next**.
- Click **NDMP Storage**, and then click **Next**.
- Enter a host name or fully qualified domain name of the NDMP server and a description, and then click **Next**.

5 Specify information for the following fields:

Port number	Specify the port number to be used for communication between the Backup Exec server and the NDMP server.
Logon account	Select the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.
Have Backup Exec use ICMP ping operations to detect the server	Select this option to enable Backup Exec to use ping to locate the NDMP server.

6 Click **Next**.

7 Review the summary, and then do one of the following:

To change the configuration	Do the following in the order listed: - Click the heading that contains the items that you want to change. - Make any changes, and then click Next until the summary appears. - Click Finish.
To configure the NDMP server	Click Finish .

8 Click the **Backup and Restore** tab, and view the NDMP server in the list of servers.

Sharing the tape drives on NDMP servers between multiple Backup Exec servers

To share the tape drives that are attached to a network-attached storage NDMP server, the NDMP feature and the Central Admin Server feature (CAS) must be installed. CAS is installed as part of the Enterprise Server feature (ESF).

You can then select which Backup Exec servers can share the tape drives. The Backup Exec server that you used to add the NDMP server is automatically selected for sharing.

Note: If you upgraded from an earlier version of Backup Exec, your existing configuration is preserved. You do not have to set up sharing for existing configurations.

To share a tape drive on NDMP servers between multiple Backup Exec servers

- 1 On the **Storage** tab, right-click the storage device that you want to share.
- 2 Click **Share**.
- 3 Check the Backup Exec servers or managed Backup Exec servers that you want to share this storage device.
- 4 Click **OK**.

See [“About the Central Admin Server feature”](#) on page 1321.

NDMP backup options for NDMP servers

When you create a backup job for a network-attached storage NDMP server, you can use the options that are appropriate for the job and the type of NDMP server.

The following backup options are available for NetApp and Fujitsu NDMP servers:

Table M-2 NDMP backup options for NetApp/Fujitsu

Item	Description
Back up Access Control Lists	Backs up NetApp Access Control Lists.
Enable file history	Enables the generation of file history data. File history is used to optimize recovery of selected backup sets, however, file history generation and processing increase the backup time. Disabling this option improves backup time. If the file history is not generated, and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Level 1 provides a differential backup that backs up new or modified files since level 0. Levels 2 through 9 back up new or modified files since the previous level backup. For example, the level 2 backup method backs up new or modified files since the level 1 backup. Level 3 backup backs up new or modified files since the level 2 backup, and so on.

The following backup options are available for IBM servers:

Table M-3 NDMP backup options for IBM

Item	Description
Back up Access Control Lists	Backs up NetApp Access Control Lists.
Enable file history	Enables the generation of file history data. File history is used to optimize recovery of selected backup sets, however, file history generation and processing increase the backup time. Disabling this option improves backup time. If the file history is not generated, and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Level 1 provides a differential backup that backs up new or modified files since level 0. Levels 2 through 9 back up new or modified files since the previous level backup. For example, the level 2 backup method backs up new or modified files since the level 1 backup. Level 3 backup backs up new or modified files since the level 2 backup, and so on.
SENDFILE	Specifies whether the jobs are allowed to use the SENDFILE environment variable. When this check box is selected, Backup Exec controls SENDFILE and sets it to the appropriate value. If the check box is not selected, the SENDFILE variable is set to 0 and is not used for the job.

The following backup options are available for EMC NDMP servers

Table M-4 NDMP backup options for EMC

Item	Description
Backup type	Determines the backup type for this backup job. The following backup types are available: ■ Dump ■ VBB
Back up with integrated checkpoints (SnapSure)	Enables Backup Exec to create a backup set that uses the EMC SnapSure feature. For more information about SnapSure, see your EMC documentation.
Enable file history	Enables the generation of file history data. File history is used to optimize the recovery of selected backup sets. File history generation and processing increase the backup time. Disabling this option improves backup time. If file history is made unavailable and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Level 1 provides a differential backup that backs up new or modified files since level 0. Levels 2 through 9 back up new or modified files since the previous level backup. For example, the level 2 backup method backs up new or modified files since the level 1 backup. Level 3 backup backs up new or modified files since the level 2 backup, and so on.

For NDMP servers other than NetApp, IBM, Fujitsu, and EMC, Backup Exec displays the appropriate options. You can change the values for most options, such as the backup level and file history option. You can change some options to a predefined value, some to any value, and some options you cannot change.

The options that appear for each NDMP server have been tested. However, while the NDMP server may support additional options, these may not be officially supported. Backup Exec does not validate the values that you enter for the options, so ensure that you enter the correct values. See the documentation for your NDMP server for information about the options that appear.

NDMP server backup selections

When you create a backup, you can select a network-attached storage NDMP server. On the NDMP server, you can select volumes or directories, or you can specify the data to include or exclude.

The following limitations apply when you select volumes or directories on NDMP servers for backup:

- You can include entire volumes for any NDMP server.
- You can include subfolders for NetApp/IBM/Fujitsu NDMP servers only.
- You cannot exclude files or directories.

See [“Backing up data”](#) on page 158.

If you want to specify what to include or exclude in the backup job, you can do the following:

- Select specific directories to include in the backup job.
- Type the names of the specific directories and files that you want to exclude from the backup job in the **Include/Exclude** dialog box.

Table M-5 What you can include and exclude for NDMP server backups

Type of NDMP server	Include	Exclude
NetApp/IBM/Fujitsu	Single or multiple directories	Directories and files, based on an exclusion pattern.
EMC	Single directory	Directories and files, but only if you select the Dump backup type.
Other	Single or multiple directories, depending on the NDMP server.	Directories and files. Refer to your NDMP server documentation for details.

See [“How to use patterns to exclude files and directories from an NDMP server backup”](#) on page 1406.

How to use patterns to exclude files and directories from an NDMP server backup

When you exclude files and directories from a backup for a NetApp/IBM/Fujitsu or EMC Celerra network-attached storage NDMP server, you must use patterns.

You should enter patterns carefully to ensure that you exclude the correct files and directories. Backup Exec does not verify the validity of exclude patterns. If you enter an invalid pattern, the pattern is ignored and therefore the files or directories are not excluded.

For details about how to use patterns, see your NDMP server documentation.

The following example shows a pattern to exclude files and directories from a backup selection for a NetApp/IBM/Fujitsu NDMP server:

Table M-6 Example pattern for NetApp/IBM/Fujitsu NDMP servers

Pattern	Example
tmp	Excludes all files and directories that have the name "tmp".
*.core	Excludes all files and directories that end with ".core".

To exclude directories for an EMC Celerra NDMP server, do not include the name of the EMC Celerra server or the name of the file system in the pattern. The names of the NDMP server and the file system are already included in the **Resource name** text box on the **Include/Exclude** dialog box. If you repeat the name of the NDMP server and the file system in the pattern, the EMC Celerra NDMP server ignores the exclusion. Type the path from the root directory to the directory that you want to exclude. Do not include an initial forward slash (/).

The following example shows a pattern to exclude directories from a backup selection for an EMC Celerra NDMP server:

Table M-7 Example pattern to exclude directories for an EMC Celerra NDMP server

Pattern	Description
test_exclusion/subdir1	Excludes only the "subdir1" directory on the file system that is listed in the Resource name text box.
test_exclusion/d*	Excludes all directories that start with the letter "d " under the directory "/test_exclusion"

The following example shows a pattern to exclude files from a backup selection for an EMC Celerra NDMP server:

Table M-8 Example pattern to exclude files for an EMC Celerra NDMP server

Pattern	Description
*.mp3	Excludes all files that end with “.mp3”.
temp	Excludes all files that have the name “temp”.

See [“NDMP server backup selections ”](#) on page 1406.

See [“About selecting data to back up ”](#) on page 171.

Supported configurations for duplicating data from NDMP servers

With the NDMP feature, you can duplicate the backup data from a network-attached storage NDMP server to storage that is attached to a Backup Exec server or to another NDMP server.

Backup Exec supports the following configurations for duplicating backup data from NDMP servers:

- Two storage devices that are attached locally to the Backup Exec server.
- Two tape drives that are attached locally to an NDMP server.
- One tape drive that is attached locally to an NDMP server and one tape drive that is attached locally to another NDMP server.
- A storage device that is attached locally to a Backup Exec server and one tape drive that is attached locally to an NDMP server.

The procedure to duplicate backup data from NDMP servers is the same as the procedure to duplicate any other type of data. For NetApp/IBM/Fujitsu NDMP servers, you must also select the logon credential for the source NDMP server.

See [“Duplicating backup sets or a job history manually”](#) on page 222.

About restoring and redirecting restore data for NDMP servers

With the NDMP feature, you can use the Restore Wizard on the **Backup and Restore** tab to restore data for a network-attached storage NDMP server. During the restore process, you can select individual files for restore if file history was enabled for the backup job. You cannot exclude files and directories when you restore to an NDMP server. Excluded directories and files are restored.

Backup Exec cannot gather sufficient file and directory information on an NDMP server restore job to accurately populate the **Backup Set Summary** and **Backup Set Information** sections of the job history. Therefore, the number of files, directories, files skipped, corrupt files, and files in use always appears as 0.

Note: Ensure that the default catalog option **Use storage-based catalogs** is selected. Otherwise, NDMP server backup sets cannot be cataloged.

About redirecting restore data for NDMP servers

You can use the NDMP feature to redirect restore data from one network-attached storage NDMP server to another NDMP server.

When you redirect NDMP server data, be aware of the following limitations:

- You cannot redirect NDMP server data to a computer that runs the Windows or Linux operating systems.
- You cannot redirect non-NDMP server data, such as NTFS or SQL data, to an NDMP server.
- The NDMP server to which you want to redirect the restored data must be from the same vendor, brand, or family as the NDMP server from which the data was backed up.

Backup Exec cannot gather sufficient file and directory information on an NDMP server backup to accurately populate the **Job Summary Information** and the **Set Detail Information** sections of the job history. Therefore, for restore and verify operations, the number of files, directories, skipped files, corrupt files, and files in use always appears as 0. Job summary and set detail information display for backup and duplicate operations.

Note: Verify operations are only supported for NetApp servers.

See [“Configuring default options for catalogs”](#) on page 255.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“NDMP server restore options”](#) on page 1409.

NDMP server restore options

With the NDMP feature, you can create a restore job for a network-attached storage NDMP server. The options that appear in the Restore Wizard vary depending on the type of NDMP server.

See [“About restoring and redirecting restore data for NDMP servers”](#) on page 1408.

Table M-9 NDMP feature restore options for NetApp/IBM/Fujitsu

Item	Description
Restore Access Control Lists	Restores NetApp Access Control Lists.
Enable Direct Access Recovery	Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backed up data stream. The NDMP server can then read the data that applies to the single file that you want to restore. This practice reduces the amount of information that is processed and significantly reduces recovery time. If you do not select this option, the restore may take significantly longer. Note: Not all vendors provide Direct Access Recovery.
Restore without writing data to disk (Verify data without doing a restore)	Tests the validity of the data that you selected for the restore job. Backup Exec does not restore the data. For NetApp/IBM NDMP servers, you should use this option to verify data instead of the Verify option on the backup definition..
Recreate the directory structure from the backup when the data is restored; otherwise, all data is restored without any directory structure	Restores the data with its original directory structure intact.

Table M-10 NDMP feature restore options for EMC

Item	Description
Enable Direct Access Recovery	Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backed up data stream. The NDMP server can then read the data that applies to the single file that you want to restore. This practice reduces the amount of information that is processed and significantly reduces recovery time. If you do not select this option, the restore may take significantly longer. Note: Not all vendors provide Direct Access Recovery.
Recreate the directory structure from the backup when the data is restored; otherwise, all data is restored without any directory structure	Restores the data with its original directory structure intact.
Restore over existing files	Overwrites the files on the restore destination that have the same name as files that are restored. Use this option only when you are sure that you want to restore an older version of a file.

For NDMP servers other than NetApp, IBM, Fujitsu, and EMC, Backup Exec displays the appropriate variables and default values for the specific type of NDMP server. You can change the values as needed. Variables that begin with the prefix "@@" are specific to Backup Exec rather than to a specific NDMP server. The options that appear for each NDMP server have been tested. However, while the NDMP server may support additional options, these may not be officially supported. Backup Exec does not validate the values for the variables that you enter, so you should ensure that you enter the values correctly. See the documentation for your NDMP server for information about the values to use.

Setting the default backup options for the NDMP feature

You can use the defaults for all NDMP server backup jobs that Backup Exec sets for the NDMP feature during installation, or you can choose your own defaults. You can also change the defaults for any specific backup job.

To set default backup options for the NDMP feature

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.
- 3 In the left pane, select **NDMP**.
- 4 Do either of the following:
 - Select the appropriate options for NetApp, IBM, and Fujitsu NDMP servers:

Back up Access Control Lists	Backs up NetApp Access Control Lists.
Enable file history	Enables the generation of file history data. File history is used to optimize recovery of selected backup sets, however, file history generation and processing increase the backup time. Disabling this option improves backup time. If the file history is not generated, and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.
■ Select the appropriate options for EMC NDMP servers	

Backup type	Determines the backup type for this backup job. The following backup types are available: ■ Dump ■ VBB
Back up with integrated checkpoints (SnapSure)	Enables Backup Exec to create a backup set that uses the EMC SnapSure feature. For more information about SnapSure, see your EMC documentation.
Enable file history	Enables the generation of file history data. File history is used to optimize the recovery of selected backup sets. File history generation and processing increase the backup time. Disabling this option improves backup time. If file history is made unavailable and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.

5 Click **Apply**.

Viewing the properties of an NDMP server

You can view the properties of a network-attached storage NDMP server that you back up.

To view the properties of an NDMP server

- 1 On the **Backup and Restore** tab, double-click the NDMP server.
- 2 In the left pane, select **Properties**.

See [“NDMP server properties”](#) on page 1414.

NDMP server properties

You can view the following properties for a network-attached storage NDMP server:

See [“Viewing the properties of an NDMP server”](#) on page 1413.

Table M-11 NDMP server properties

Item	Description
Name	Indicates the name of the NDMP server.
Description	Shows the user-defined description of the NDMP server.
Logon account	Indicates the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.

Viewing storage properties for an NDMP server

You can view the storage properties of a network-attached storage NDMP server.

To view storage properties for an NDMP server

- 1 On the **Storage** tab, double-click the NDMP server.
- 2 In the left pane, select **Properties**.

See [“Storage properties for an NDMP server”](#) on page 1414.

Storage properties for an NDMP server

You can view the following storage properties for a network-attached storage NDMP server.

See [“Viewing storage properties for an NDMP server”](#) on page 1414.

Table M-12 NDMP server storage properties

Item	Description
Server name	Indicates the name of the NDMP server.
Description	Shows the user-defined description of the server.

Table M-12 NDMP server storage properties (continued)

Item	Description
State	Indicates the status of the NDMP server storage. See “Backup Exec server and storage device states” on page 584.
Port	Lists the port that is used for communications between the Backup Exec server and the NDMP server.
Use ICMP ping operations to detect the server	Indicates whether ICMP ping is enabled. ICMP ping enables Backup Exec to use ping to locate the NDMP server.
Logon account	Indicates the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.
Host ID	Displays the identifier number that the NDMP server generates.
System version	Indicates the software version that is installed on the NDMP server.

About support for NetApp's Data ONTAP

Backup Exec supports Node-scope NDMP mode and VServer Aware NDMP mode for NetApp's data ONTAP. In the Node-scope mode, the node names are used as the NDMP host name.

In VServer mode, Backup exec supports Cluster Aware Backup (CAB) which is an NDMP v4 protocol extension. This extension allows you to back up a Virtual Server (VServer) or a Storage Virtual Machine (SVM) and enables the NDMP server to establish a data connection on a node that owns a volume. This also enables Backup Exec to determine if volumes and tape devices are located on the same node in a cluster.

Backup Exec File Servers

This appendix includes the following topics:

- [About AWS FSx for Windows File Server](#)
- [About Azure Files](#)
- [Pre-requisites for AWS FSx and Azure Files](#)
- [Notes for AWS FSx and Azure Files](#)
- [Recommendation for AWS FSx and Azure Files](#)
- [Best practices for AWS FSx and Azure Files](#)
- [Adding AWS FSx or Azure Files to Backup Exec](#)
- [Backing up AWS FSx or Azure Files](#)
- [Restoring AWS FSx or Azure Files](#)
- [Limitation of Azure Files](#)

About AWS FSx for Windows File Server

Amazon FSx for Windows File Server or AWS FSx provides fully managed, highly reliable file storage that is accessible over the industry-standard Service Message Block (SMB) protocol. It is built on Windows Server with Microsoft Active Directory (AD) integration.

About Azure Files

Azure file shares are serverless, deploying for production scenarios does not require managing a file server or NAS device. Azure file shares are deployed into

storage accounts, which are top-level objects that represent a shared pool of storage. Azure file shares get created at storage account level.

Pre-requisites for AWS FSx and Azure Files

For AWS FSx, ensure that Backup Exec is configured in the Active Directory environment. Ensure that all the pre-requisites to access AWS FSx from your network are in place. For more information about AWS FSx prerequisites, refer to the following link:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/walkthrough01-prereqs.html>

<https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html>

For Azure file shares, ensure that all the pre-requisites to access Azure file shares from your network are in place. For more information about Azure file share pre-requisites, refer to the following link:

<https://docs.microsoft.com/en-us/azure/storage/files/storage-files-networking-overview>

Notes for AWS FSx and Azure Files

You should review the following notes before you use for AWS FSx and Azure file shares:

- Backup Exec can be deployed on-premises or in the cloud, to protect Azure file shares and AWS FSx shares.
- Backup throughput rate depends on the internet bandwidth available between the on-premises network and the AWS or Azure infrastructure.
- Ensure AWS FSx shares or Azure file servers are accessible and you can access them as SMB shares from the Backup Exec server. You must have full rights for AWS FSx shares and Azure file servers.

Recommendation for AWS FSx and Azure Files

Recommendations to help you use AWS FSx and Azure file shares.

- For redirected restore, it is recommended to use AWS FSx shares, that is configured in the same domain.
- For redirected restore, it is recommended to use the Azure file shares destination from the same Storage account.

Best practices for AWS FSx and Azure Files

Review the following best practice before you use AWS FSx and Azure file shares:

- Configure AWS FSx shares and Azure file shares backup job over cloud connectors as a storage device to help reduce cost.

Adding AWS FSx or Azure Files to Backup Exec

You can use the **Add Server Wizard** in Backup Exec to add DNS name of AWS FSx file system (AWS FSx share) or endpoint of the Azure Storage account. After you add the DNS of the file system and endpoint, you can back up the contents.

To add AWS FSx shares or Azure Files

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **File server, NDMP data server, Azure file share, or AWS FSx for Windows File Server**, and then click **Next**.
- 3 Select **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Enter the DNS name of AWS FSx share or endpoint of the Azure Storage account and then click **Next**.
- 5 In the **Logon Account** field, select one of the following:

For Amazon FSx Enter the logon account used while creating AWS FSx share on the AWS portal.

For Azure Enter the following:

```
Azure\<Storage_account_name>  
<Access_key_of_storage_account>
```

- 6 Review the summary information, and then click **Finish**.

Backing up AWS FSx or Azure Files

You can back up contents from AWS FSx and Azure file shares from Backup Exec. AWS FSx shares and Azure file shares support all the backup methods, such as full, incremental, and differential backups.

While configuring backup job definition, click **Test/Edit credentials** option and select Azure Storage account's credentials. The backup job will fail if the credentials are incorrect.

To back up AWS FSx or Azure Files

- 1 On the **Backup and Restore** tab, right-click the DNS name of AWS FSx share or the Azure Storage account's endpoint that you want to back up.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, select the Azure file shares or AWS FSx share contents that you want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 Set any additional options for this job.
- 8 Click **OK**.

See [“Backing up data”](#) on page 158.

Restoring AWS FSx or Azure Files

You can perform a restore to the original location or you can redirect the restore to a new location. For redirected restore, you need to recover files and folders to another computer that has a Remote Agent for Windows installed.

To restore AWS FSx or Azure Files

- 1 On the **Backup and Restore** tab, right-click the DNS name of AWS FSx share or the Azure Storage account's endpoint for which you want to restore data, and then click **Restore**.
- 2 Select **Files, folders, or volumes**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“Restoring file system data”](#) on page 244.

Limitation of Azure Files

If you edit an existing file directly in the Azure portal, the modified time of the file in Windows explorer does not change, even though the file content is updated. As a result, when an incremental backup runs, the changes in the file are not backed up.

It is a known limitation from Azure. For more information, refer to the Azure File Sync section in the Microsoft Azure documentation.

<https://docs.microsoft.com/en-us/azure/storage/files/storage-files-faq>



Backup Exec Agent for Linux and Unix

This appendix includes the following topics:

- [About the Agent for Linux and Unix](#)
- [About open files and the Agent for Linux](#)
- [Requirements for the Agent for Linux and Unix](#)
- [About installing the Agent for Linux and Unix](#)
- [About establishing trust for a remote Linux and Unix computer in the Backup Exec list of servers](#)
- [Adding additional Backup Exec servers to which the Agent for Linux and Unix can publish information](#)
- [About configuring the Agent for Linux and Unix](#)
- [Excluding files and directories from all backup jobs for Linux and Unix computers](#)
- [Editing configuration options for Linux and Unix computers](#)
- [About backing up a Linux and Unix computer by using the Agent for Linux and Unix](#)
- [About backing up Linux and Unix shares without using the Agent for Linux and Unix](#)
- [About restoring data to Linux and Unix computers](#)
- [Editing the default backup job options for Linux and Unix computers](#)

- [Uninstalling the Agent for Linux and Unix](#)
- [Starting the Agent for Linux and Unix daemon](#)
- [Stopping the Agent for Linux and Unix daemon](#)
- [Troubleshooting the Agent for Linux and Unix](#)

About the Agent for Linux and Unix

The Backup Exec Agent for Linux and Unix (Linux and Unix Agent) is installed as a separate add-on component. The Linux and Unix Agent enables network administrators to perform backup and restore operations on Linux and Unix servers that are connected to the network. The Linux and Unix Agent must be installed on the Linux and Unix servers before you can perform backup or restore operations.

For information about the best practices to use Backup Exec Agent for Linux and Unix (Linux and Unix Agent), refer to *Backup Exec Best Practices*.

See [“About open files and the Agent for Linux”](#) on page 1422.

See [“Requirements for the Agent for Linux and Unix”](#) on page 1423.

See [“About installing the Agent for Linux and Unix”](#) on page 1423.

About open files and the Agent for Linux

The Agent for Linux uses advanced open file and image technologies that are designed to alleviate the issues that are sometimes encountered during backup operations, such as backing up open files.

After you make file and folder selections and submit the job for backup, the Linux Agent automatically makes a snapshot of the volume or volumes. Making a snapshot of a volume provides a point-in-time record of the data. When the Linux Agent creates a snapshot, it uses snapshot technologies to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. During the backup, files can be open and data can be changed.

The Linux Agent supports Simple, Logical Volume Manager (LVM), and RAID volume configurations.

See [“Requirements for the Agent for Linux and Unix”](#) on page 1423.

Requirements for the Agent for Linux and Unix

The following items are required to install the Agent for Linux and Unix (Linux and Unix Agent):

- The Backup Exec server must have TCP/IP installed.
- The Linux and Unix server must have Perl 5.8.8 or later installed.
- You must have a root logon account on the Linux and Unix servers.
- You must have the Backup Exec installation media.
- You must enter a license for the Linux and Unix Agent on the Backup Exec server.

Note: Some versions of Linux may require that you install the libstdc++.so.5 package.

See [“Troubleshooting the Agent for Linux and Unix”](#) on page 1450.

It is recommended that you use the Secure Shell (SSH) protocol when you push-install the Linux and Unix Agent from one Linux and Unix server to another Linux and Unix server. You must enable SSH before you push-install the Linux and Unix Agent.

Backup Exec automatically installs the Remote Media Agent for Linux when it installs the Agent for Linux on a Linux server. However, you must enter a separate license for the Remote Media Agent for Linux before it is available for use. It applies specifically for Linux.

See [“About the Remote Media Agent for Linux ”](#) on page 1456.

You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

See [“About installing the Agent for Linux and Unix”](#) on page 1423.

About installing the Agent for Linux and Unix

Use the Backup Exec installation media to do the following:

- Install the Agent for Linux and Unix (Linux and Unix Agent) on a local Linux and Unix server.
- Push-install the Linux Agent from one Linux server to other remote Linux servers.

If you push-install the Linux Agent, the RSH (Remote Shell) is used by default. It is recommended that you use SSH (Secure Shell) instead. To use SSH, you must enable the SSH before you install the Linux Agent. Refer to your operating system documentation for more information about SSH. It applies specifically for Linux.

Before you install the Linux and Unix Agent, you should review the requirements:

See [“Requirements for the Agent for Linux and Unix”](#) on page 1423.

When you install the Linux and Unix Agent, Backup Exec creates the beoper group and adds root as a member. The beoper group contains the names of the users who have permission to back up and restore the Linux and Unix servers. However, if Backup Exec detects an NIS server during the Linux and Unix Agent installation, then the beoper group is not created. You must create the beoper group manually on the Linux and Unix servers on which you want to install the Linux and Unix Agent.

When the installation is complete, Backup Exec saves the install log file to the following location on the server on which the Linux and Unix Agent is installed:

```
/var/tmp/vxif/installralus<summary file number>/installralus.log
```

See [“Installing the Agent for Linux and Unix ”](#) on page 1424.

Installing the Agent for Linux and Unix

You can install the Agent for Linux and Unix (Linux and Unix Agent) on a local Linux and Unix server.

You can then push-install the Linux Agent from the local Linux server to one or more remote Linux servers. It applies specifically for Linux.

See [“About installing the Agent for Linux and Unix”](#) on page 1423.

Note: You must unzip the RALUS_RMALS_<version number>.gz file on a Linux system. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Agent for Linux and Unix

- 1 At a Linux and Unix server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install the Linux and Unix Agent.

- 3 Navigate to the following directory on the installation media.
<Unix>
- 4 Copy the **RALUS_RMALS_<version number>.gz** file in this directory to a directory on the local computer.
- 5 Unzip the file.
For example:

```
gunzip RALUS_RMALS_<version number>.gz
```
- 6 Untar the file.
For example:

```
tar -xf RALUS_RMALS_<version number>.tar
```
- 7 Do one of the following:

To install the Linux and Unix Agent on the local Linux and Unix server	Start the installralus script. For example: <code>./installralus</code>
To install the Linux and Unix Agent from the local Linux and Unix server to one or more remote Linux and Unix server	Do the following in the order listed: ■ Start the installralus script using the -SSH switch. For example: <code>./installralus -usessh</code> ■ Type the name, IP address, or fully qualified domain name of a Linux and Unix server. Note: To install the agent to multiple remote Linux and Unix servers, leave a space between each identifier.
- 8 After the installer checks for a valid Linux or Unix operating system during the initial system check, press **Enter**.
- 9 Review the package installation summary, and then press **Enter**.
- 10 After the system installation requirements check completes, press **Enter**.
- 11 Start the prerequisites check by pressing **Enter**.
- 12 Start the NIS server scan by pressing **Enter**.
- 13 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected	The Linux and Unix Agent installer cannot create the beoper group. You must create it manually after the Linux and Unix Agent installation is complete. Continue with the next step.
If an NIS server is not detected	Use the installer to create the beoper group. Do the following in the order listed: ■ To let the installer create the beoper group, type y. ■ To select the next available Group ID, type n. ■ To add the root user account to the beoper group, type y. ■ Continue with the next step.

- 14 Start the installation by pressing **Enter**.
- 15 After installation completes, press **Enter** to start the post-installation configurations and installation of SymSnap drivers.

Note: It applies specifically for Linux.

- 16 Press **Y** to automatically start the Beremote service; otherwise, press **N** to start the service later.
- 17 After the configuration process completes, press **Enter** to save the installation log to the following file:
`/var/tmp/vxif/installralussummary file number/installralus.log`
- 18 If the Linux and Unix Agent installer did not create a beoper group, you must create it.

See [“Creating the Backup Exec operators \(beoper\) group manually”](#) on page 1427.
- 19 Configure the Agent for Linux and Unix as appropriate.

See [“About configuring the Agent for Linux and Unix”](#) on page 1430.
- 20 If the Beremote service is not running, start the Agent for Linux and Unix daemon.

See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

About the Backup Exec operators (beoper) group for the Agent for Linux and Unix

The beoper group contains the names of the users who have permission to back up and restore the Linux and Unix servers.

When you install the Agent for Linux and Unix (Linux and Unix Agent), Backup Exec creates the beoper group and adds root as a member. Any Linux user that you add to the beoper group gets the necessary permissions to back up and restore the servers.

However, if an NIS server is detected during the Linux and Unix Agent installation, Backup Exec cannot create the beoper group. You must create the beoper group manually on the Linux and Unix servers on which you want to install the Linux and Unix Agent. You must create the beoper group before you start backup and restore operations. Otherwise, connections fail between the Linux and Unix servers and the Backup Exec server.

Before the members of the beoper group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators \(beoper\) group manually”](#) on page 1427.

See [“Backup Exec logon accounts”](#) on page 759.

Creating the Backup Exec operators (beoper) group manually

You must create a beoper group on each server on which you want to install the Agent for Linux and Unix (Linux and Unix Agent).

See [“About the Backup Exec operators \(beoper\) group for the Agent for Linux and Unix”](#) on page 1427.

Note: Ensure that you understand how to set security for groups on Linux and Unix servers before you assign a Group ID for the beoper group.

Table O-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux and Unix server on which you want to install the Linux and Unix Agent. If the Linux and Unix server is in an NIS domain, navigate to the NIS domain's group file.	Refer to the NIS documentation for information on how to add a group to an NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system's documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux and Unix server.	See the operating system's documentation for more information about how to add users to a group.
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See “Backup Exec logon accounts” on page 759.

About establishing trust for a remote Linux and Unix computer in the Backup Exec list of servers

When you connect to a Linux and Unix computer from the Backup Exec server, you must establish trust between the Backup Exec server and the remote Linux and Unix computer. You must also establish trust if you want to configure a remote Linux and Unix computer to perform client-side deduplication.

Note: Client side deduplication is supported only for Linux.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 953.

See [“Establishing trust and adding a remote Linux and Unix computer to the Backup Exec list of servers”](#) on page 1429.

Establishing trust and adding a remote Linux and Unix computer to the Backup Exec list of servers

You can add one or more remote Linux and Unix computers to the list of servers that appear on the **Backup and Restore** tab. When you add remote Linux and Unix computers, you must establish a trust between the Backup Exec server and the remote Linux and Unix computers to ensure secure communication.

To establish trust and add a remote Linux and Unix computer to the Backup Exec list of servers

- 1 On the **Backup and Restore** tab, in the **Servers** group, click **Add**.
- 2 Click **Unix computer**.

Note: Add server displays **Unix computer** for Linux and Unix.

- 3 Follow the on-screen prompts.

See [“Adding additional Backup Exec servers to which the Agent for Linux and Unix can publish information”](#) on page 1429.

See [“About configuring the Agent for Linux and Unix”](#) on page 1430.

Adding additional Backup Exec servers to which the Agent for Linux and Unix can publish information

You can specify additional Backup Exec servers to which the Agent for Linux and Unix (Linux and Unix Agent) can publish information.

Each Backup Exec server to which the Linux and Unix Agent publishes information appears in the Backup Exec **Servers** list.

To add additional Backup Exec servers to which the Agent for Linux and Unix can publish information

- 1 Use a text editor to open the following file:

```
/etc/VRTSralus/ralus.cfg
```

- 2 Add the following string:

```
Software\Veritas\Backup Exec For Windows\Backup  
Exec\Engine\Agents\Agent Directory List unique identifier number = IP  
address or DNS name of Backup Exec server
```

- 3 Save and close the file.
- 4 Move to the Backup Exec server to which the Linux and Unix Agent is publishing itself and add the Linux and Unix server to the **Servers** list.

See [“Adding servers that you want to back up to the list of servers on the Backup and Restore tab”](#) on page 151.

About configuring the Agent for Linux and Unix

Backup Exec creates a file named `ralus.cfg` on each Linux and Unix server on which the Agent for Linux and Unix (Linux and Unix Agent) is installed. You can edit the strings, identifiers, and variables in this file to add or edit options for the Linux and Unix Agent.

Options that you can edit in the `ralus.cfg` file include the following:

- The port to which the Linux and Unix Agent must send publishing messages.
- The logging level for Oracle database operations that use the Backup Exec Linux and Unix Agent Utility, and for NDMP information.
- The settings to allow the Linux and Unix Agent to publish to one or more Backup Exec servers.
- The files and directories on Linux and Unix servers that you want to exclude from backups.

The `ralus.cfg` file format contains three components. The first component (A) in the following example is a required string.

The second component (B) is a unique identifier followed by an equal sign (=). A unique identifier can consist of sequential numbers, letters, or alpha-numeric characters. For example, 1, 2, 3 or A, B, C. You can also use AA, BB, CC, or A1, A2, B1, B2.

The third component of the `ralus.cfg` format is the NetBIOS name, fully qualified domain name, or IP address of the Backup Exec server.

The `ralus.cfg` includes a registry key that works with the Linux Agent's open file technology. It applies specifically for Linux. The name of the key is `DisableOFO` and appears in the `ralus.cfg` file in the following form:

```
Software\Veritas\Backup Exec for Windows\Backup  
Exec\Engine\RALUS\DisableOFO=0
```

By default, the `DisableOFO` key is set to 0, meaning that the Linux Agent is active, letting the Linux Agent back up the open files that it encounters. However, you

can disable the open file technology by changing the value of the key to "1", and then restarting the Linux Agent daemon.

Figure O-1 Example of the `ralus.cfg` file

A

B

C

Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 1=srv.mycompany.com
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 2=datasrv
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 3=66.35.250.151

A = Required string
B = Required and unique identifier (the order or appearance is irrelevant)
C = File or directory to be excluded

- See [“Editing configuration options for Linux and Unix computers”](#) on page 1432.
- See [“Configuration options for Linux and Unix computers”](#) on page 1432.
- See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.
- See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

Excluding files and directories from all backup jobs for Linux and Unix computers

You can exclude specific files and directories on the Linux and Unix computers from all backup jobs. Edit the `ralus.cfg` file to specify the excluded files.

The following is an example of strings in the `ralus.cfg` file that excludes files and directories from all backup jobs.

Figure O-2 Example of file and directory exclusions in the `ralus.cfg` format

A

B

C

Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude1=/dev/*.*
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude2=/proc/*.*
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude3=/mnt/nss/pools/
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude4=/mnt/nss/.pools/

A = Required string
B = Required and unique identifier (the order or appearance is irrelevant)
C = File or directory to be excluded

To exclude files and directories for specific backup jobs, specify the exclusions in the backup job properties.

See [“Editing configuration options for Linux and Unix computers”](#) on page 1432.

Editing configuration options for Linux and Unix computers

You can edit configuration options for the Agent for Linux and Unix.

See [“About configuring the Agent for Linux and Unix”](#) on page 1430.

To edit configuration options for Linux and Unix computers

- 1 Use a text editor to open the following file:

```
/etc/VRTSralus/ralus.cfg
```

- 2 Change the appropriate string in the file.

See [“Configuration options for Linux and Unix computers”](#) on page 1432.

Configuration options for Linux and Unix computers

You can edit options to configure the Agent for Linux and Unix (Linux and Unix Agent).

See [“Editing configuration options for Linux and Unix computers”](#) on page 1432.

Table O-2 Configuration options for Linux and Unix computers

String and default values	Description
Software\Veritas\Backup Exec For Windows\Agent Browser\TcpIp\AdvertisementPort=6101	Lists the port to which the Linux and Unix Agent must send publish and purge messages.
Software\Veritas\Backup Exec for Windows\Backup Exec\Debug\AgentConfig=0	Enables logging for the Linux and Unix Agent utility that Oracle operations use. Values include the following: ■ 0 Logging is not enabled. ■ 1 Logging is enabled. Backup Exec automatically generates the log file.

Table O-2 Configuration options for Linux and Unix computers (*continued*)

String and default values	Description
Software\Veritas\Backup Exec for Windows\Backup Exec\Debug\VBXBSAlevel=0	Enables logging for the Linux and Unix Agent for Oracle operations. Values include the following: ■ 0 Logging is not enabled. ■ 5 Normal logging is enabled. ■ 6 Advanced logging is enabled. Large log files may be created.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\Agents\Advertise All=1	Enables the Linux and Unix Agent to publish information to all of the Backup Exec servers that are listed in the \Agents\Agent Directory List strings. Values include the following: ■ 1 The Linux and Unix Agent publishes information to every Backup Exec server in the Agent Directory List. ■ 0 The Linux and Unix Agent publishes information to the first Backup Exec server in the Agent Directory List. If the attempt is successful, the Linux and Unix Agent does not publish information to any other Backup Exec servers. If the attempt is not successful, the Linux and Unix Agent attempts to publish information to the next Backup Exec server in the list. Attempts continue until the Linux and Unix Agent reaches the end of the list.

Table O-2 Configuration options for Linux and Unix computers *(continued)*

String and default values	Description
Software\Veritas\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertise Now=0	Enables the Linux and Unix Agent to start a new publishing cycle after you add or edit any settings in the ralus.cfg file. Values include the following: ■ 0 The Linux and Unix Agent publishes information according to its regular cycle, which is set in the string \Agents\Advertising Interval Minutes. Any changes to the ralus.cfg file take effect when a new publishing cycle begins. ■ 1 The Linux and Unix Agent starts a new publishing cycle. Any changes to the ralus.cfg file take effect immediately. If the Backup Exec server does not receive the publishing information, the Linux and Unix Agent makes 10 more attempts. Each attempt to publish information to the Backup Exec server is one minute apart. If the information is not sent at the end of the 10 attempts, the Linux and Unix Agent skips that Backup Exec server until the next publishing cycle. The publishing cycle is the number of minutes set in the string \Agents\Advertising Interval Minutes.

Table O-2 Configuration options for Linux and Unix computers (*continued*)

String and default values	Description
Software\Veritas\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertisement Purge=0	Lets the Linux and Unix Agent send a purge message to all of the Backup Exec servers in the string \Agents\Advertisement Purge. When a Backup Exec server receives a purge message, it removes the Linux and Unix Agent from Backup Exec's list of available servers. The Linux and Unix Agent continues to function. Values include the following: ■ 0 Do not purge the Linux and Unix Agent from any Backup Exec servers that are listed in the \Agents\Advertisement Purge string. ■ 1 Purge the Linux and Unix Agent from one or more Backup Exec servers in the \Agents\Advertisement Purge string.
	Enables the Linux and Unix Agent to publish to Backup Exec servers. Values include the following: ■ 0 The Linux and Unix Agent attempts to publish information to the Backup Exec servers that are listed in the string \Agents\Agent Directory List. ■ 1 The Linux and Unix Agent does not publish information to Backup Exec servers.
Software\Veritas\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertising Interval Minutes=240	Sets the number of minutes that the Linux and Unix Agent must wait between publishing cycles. The default number of minutes is 240. The range of minutes is from 1 minute to 720 minutes.

Table O-2 Configuration options for Linux and Unix computers (*continued*)

String and default values	Description
Software\Veritas\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List_1=<Backup Exec server name>	Displays the list of NetBIOS names, fully qualified domain names, or IP addresses to which the Linux and Unix Agent publishes information. The Backup Exec server from which the Linux and Unix Agent is push-installed is added to the Agent Directory List by default.
Software\Veritas\Backup Exec For Windows\Backup Exec\Engine\Agents\Auto Discovery Enabled=1	Adds a Backup Exec server to the string \Agents\Agent Directory List if the Backup Exec server performs a backup job with which the Linux and Unix Agent is associated. Values include the following: ■ 1 Adds the Backup Exec server that performs the backup job to the Agent Directory List. The Linux and Unix Agent can publish information to the Backup Exec server. ■ 0 The Backup Exec server that performs the backup job is not added to the Agent Directory List.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\Logging\RANT NDMP Debug Level=0	Displays the level of verbosity for logging NDMP information for the Linux and Unix Agent. Values include the following: ■ 0 Logs only the NDMP errors. ■ 1 Logs the NDMP errors and warnings. ■ 2 Logs the NDMP errors, warnings, and message information that is sent between the remote computer and the Backup Exec server.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\Encoder=	Displays the encoder that you can add if the default encoder incorrectly displays characters on the user interface.

Table O-2 Configuration options for Linux and Unix computers (*continued*)

String and default values	Description
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\ShowTSAFS=	Lets you perform a Target Service Agent file system (TSAFS) backup for applications on Novell Open Enterprise Services. By default, this option is not enabled. The Linux and Unix Agent backs up all file systems using the Root object. If ShowTSAFS is enabled, the Novell Open Enterprise Services resource appears in the backup selection list. If you select the whole computer for backup, then redundant backups are performed. It is recommended that you do not enable this option. Values include the following: ■ Blank or 0 The file system TSA does not appear for backup selection. ■ 1 The file system TSA resource appears for backup selection.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemExclude1=	Lists the files that you want to exclude from all Linux and Unix Agent backup jobs. See “ Excluding files and directories from all backup jobs for Linux and Unix computers” on page 1431.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemFSTypeExclude1	Lists the type of file system that you want to exclude from the Linux and Unix Agent backup.

Table O-2 Configuration options for Linux and Unix computers *(continued)*

String and default values	Description
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0	Lets you use the Remote Media Agent for Linux to back up the Linux server on which it is installed. By default, this option is not enabled. If you install the Remote Media Agent for Linux to an unsupported version of Linux, the Remote Media Agent for Linux is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Agent for Linux component. This component is installed with the Remote Media Agent for Linux. You must change the value of this string to 1 to use the Agent for Linux component. Values include the following: ■ 0 You can create backup, restore, and utility jobs on the Backup Exec server that run on the Linux server's storage devices. ■ 1 You can only use the Agent for Linux component to back up the Linux server on which it is installed. See “Troubleshooting the Agent for Linux and Unix” on page 1450.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\DisableOFO=0	Lets you disable or enable the Linux and Unix Agent's open file technology. By default, the DisableOFO key is set to 0, meaning that the Linux and Unix Agent is active, letting the Linux and Unix Agent back up the open files that it encounters. However, you can disable the open file technology by changing the default value of the key to 1, and then restarting the Linux and Unix Agent daemon.

Table O-2 Configuration options for Linux and Unix computers *(continued)*

String and default values	Description
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\AOFOL\CacheFileMountPoint =	The alternate cache file location for Advanced Open File feature (AOFO). This location should be the mount point. By default, the snapshot cache file is created on the Volume mount point. If you specify the location, the Linux Agent uses this location to create the snapshot cache file.
Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RALUS\AOFOL\CacheFileSize =	The alternate cache file size. The size should be at least 10% of the volume size that you plan to back up and the size value must be multiples of 4-KB blocks. The Linux Agent accepts the alternate cache file location and alternate cache file size when values for both are correctly configured. Note: This key is applicable only for the Linux agent.

About backing up a Linux and Unix computer by using the Agent for Linux and Unix

The following backup methods appear when you use the Agent for Linux and Unix (Linux and Unix Agent) to back up data:

- Full - Using modified time
- Differential - Using modified time
- Incremental - Using modified time

See [“Linux and Unix backup options”](#) on page 1439.

See [“Adding a stage to a backup definition”](#) on page 220.

See [“Editing a stage”](#) on page 221.

Linux and Unix backup options

The following Agent for Linux and Unix (Linux and Unix Agent) options are available when you back up Linux and Unix computers.

See [“About backing up a Linux and Unix computer by using the Agent for Linux and Unix”](#) on page 1439.

Table O-3 Backup job options for Linux and Unix computers

Item	Description
Preserve file and directory timestamps during backups	Prevents the Linux and Unix Agent from changing an object's attributes when a backup occurs. An object is a file or a directory. This option is not selected by default. During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime. An object's ctime is the time when an object's attributes, such as permissions and timestamps, are modified. If the Linux and Unix Agent does not change the attributes after a backup, the object's ctime is not changed. This option does not affect the object attributes that are set during restore operations.
Follow local mount points	Lets Backup Exec follow local mount points when it backs up data. This option is enabled by default. For more information on local mount points, see your operating system's documentation.

Table O-3 Backup job options for Linux and Unix computers *(continued)*

Item	Description
Follow remote mount points	Lets Backup Exec follow remote mount points when it backs up data. This option is not selected by default. When you use this option, the following limitations apply: ■ The data that is mounted must reside on a computer type that Backup Exec supports. You can find a list of supported operating systems, platforms, and applications in the Backup Exec Software Compatibility List. ■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. For more information on remote mount points, see your operating system's documentation.
Lock remote files to prevent applications from modifying them during backups	Lets the Linux and Unix Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup.

About backing up Linux ad Unix shares without using the Agent for Linux and Unix

The following backup methods appear when you do not use the Agent for Linux and Unix (Linux and Unix Agent) to back up data:

- Full - Using modified time and catalog-based backups
- Differential - Using modified time and catalog-based backups
- Incremental - Using modified time and catalog-based backups

This backup method uses the Media Agent to backup the Linux and Unix shares.

About restoring data to Linux and Unix computers

You can specify restore job options to restore Linux and Unix computers.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

See [“Restore job options for Linux and Unix computers”](#) on page 1442.

Restore job options for Linux and Unix computers

See [“About restoring data to Linux and Unix computers”](#) on page 1441.

Table O-4 Restore job options for Linux and Unix computers

Item	Description
Lock remote files if the mount points have necessary permissions	Lets Backup Exec have exclusive access to the files on the remote computers that are connected through the Network File System (NFS). This option is enabled by default.
Restore DIB set	Restores the Directory Information Base (DIB), also known as the Novell directory services (NDS) database.
Activate DIB after verify	Lets Backup Exec rename the database from .RST to .NDS after the verification process completes successfully. If the verify operation fails, the .RST file is deleted and the original .NDS file is kept intact. If you do not select this option, after the database is restored, the .RST file is available for you to perform manual activation or manual disaster recovery.
Open database when finished	Lets Backup Exec open the database after the restore completes. If you want to perform maintenance tasks before the database opens, do not select this option.
Verify database after restore	Lets Backup Exec verify the database after the restore completes.
Roll forward log directory	Displays the location of the roll forward log directory.

Editing the default backup job options for Linux and Unix computers

You can edit the existing default options for all backup and restore jobs for Linux systems.

To edit default backup job options for Linux and Unix systems

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select either **Back Up to Disk** or **Back Up to Tape**, and then select **Unix**.
- 3 Set the appropriate options.
See [“Default backup job options for Linux and Unix computers”](#) on page 1443.
- 4 Click **OK**.

Default backup job options for Linux and Unix computers

You can set default backup job properties for all jobs on Linux and Unix computers.

See [“Editing the default backup job options for Linux and Unix computers”](#) on page 1443.

You can find a list of supported operating systems, platforms, and applications in the Backup Exec Software Compatibility List.

Table O-5 Default backup job options for Linux and Unix computers

Item	Description
Preserve file and directory timestamps during backups	Prevents the Linux and Unix Agent from changing an object's attributes when a backup occurs. An object is a file or a directory. This option is not selected by default. During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime. An object's ctime is the time when an object's attributes, such as permissions or timestamps, have been modified. If the Linux and Unix Agent does not change the attributes after a backup, the object's ctime is not changed. This option does not affect the object attributes that are set during restore operations.
Follow local mount points	Lets Backup Exec follow local mount points when it backs up data. This option is enabled by default. For more information on local mount points, see your operating system's documentation.
Follow remote mount points	Lets Backup Exec follow remote mount points when it backs up data. This option is not selected by default. When you use this option, the following limitations apply: ■ The data that is mounted must reside on an operating system that Backup Exec supports. ■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. For more information on remote mount points, see your operating system's documentation.
Lock remote files to prevent applications from modifying them during backups	Lets the Linux and Unix Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup or restore job.

Uninstalling the Agent for Linux and Unix

An automated uninstall process for the Agent for Linux and Unix (Linux and Unix Agent) is available on the Backup Exec installation media.

`/opt/VRTS/install/logs/uninstallralus<summary file number>.summary`

To uninstall the Agent for Linux and Unix

- 1 On the Linux and Unix server, place the Backup Exec installation media in the appropriate device.
- 2 Log on as root to the server from which you want to uninstall the Linux and Unix Agent.
- 3 Navigate to the following directory on the Backup Exec installation media:
<Unix>
- 4 Start the **uninstallralus** script.

For example:

```
./uninstallralus
```

- 5 To uninstall the Linux and Unix Agent from one or multiple server, type the name, IP address, or fully qualified domain name of a Linux and Unix server.

Note: For multiple servers, leave a space between each identifier.

- 6 Press **Enter**.
- 7 After the Linux and Unix Agent package check completes successfully, press **Enter**.
- 8 When you are prompted to uninstall the RALUS packages, press **Enter**.
- 9 When you are prompted to uninstall the SymSnap driver, press **Enter**.

Note: It applies specifically for Linux.

- 10 To save the uninstall summary to the following location on the Linux and Unix server, press **Enter**:

`/opt/VRTS/install/logs/uninstallralus<summary file number>.summary`

See [“Installing the Agent for Linux and Unix ”](#) on page 1424.

Manually uninstalling the Agent for Linux and Unix

You can manually uninstall the Agent for Linux and Unix (Linux and Unix Agent).

To manually uninstall the Agent for Linux and Unix

- 1 Use a terminal session to connect to the Linux and Unix server as the root user.

- 2 Change to the following directory:

`/opt/VRTSralus/bin`

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Delete the following line if it is found in the `/etc/inittab` file:

`/opt/VRTSralus/bin/VRTSralus.init`

For example:

```
rm -r /opt/VRTSralus/bin/VRTSralus.init
```

- 4 Copy the `RALUS_RMALS_<version number>.gz` file in this directory to a directory on the local computer.

- 5 Unzip the file using the following command:

```
gunzip RALUS_RMALS_<version number>.gz
```

- 6 Untar the file using the following command:

```
tar -xf RALUS_RMALS_<version number>.tar
```

- 7 Stop the Linux and Unix Agent daemon.

See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.

- 8 Remove the Linux and Unix Agent package from the Linux and Unix server.

For example:

Debian GNU/Linux, Ubuntu	<code>dpkg -r VRTSralus</code>
Linux	<code>rpm -e VRTSralus</code>
Solaris	<code>pkgrm VRTSralus</code>
AIX	<code>installp -u VRTSralus</code>

9 Change back to the root directory.

For example:

```
cd /
```

10 Remove the following files:

```
/etc/VRTSralus
```

```
/opt/VRTSralus
```

```
/var/VRTSralus
```

For example:

```
rm -r /etc/VRTSralus /opt/VRTSralus /van/VRTSralus
```

11 Type **y** if you are prompted to descend into the directories.

12 Type **y** if you are prompted to delete a directory.

13 Remove run-time scripts if they are present.

See [“Run-time scripts to remove when manually uninstalling the Agent for Linux and Unix”](#) on page 1447.

Run-time scripts to remove when manually uninstalling the Agent for Linux and Unix

When you manually uninstall the Agent for Linux and Unix (Linux and Unix Agent), remove the following run-time scripts if they are present.

Table O-6 Run-time scripts to remove when manually uninstalling the Linux and Unix Agent

Operating system	Run-time scripts to remove
Debian, Ubuntu	/etc/rc5.d/S95VRTSralus.init /etc/rc3.d/S95VRTSralus.init /etc/rc2.d/S95VRTSralus.init /etc/init.d/VRTSralus.init For example: rm /etc/rc5.d/S95VRTSralus.init

Table O-6 Run-time scripts to remove when manually uninstalling the Linux and Unix Agent (*continued*)

Operating system	Run-time scripts to remove
Red Hat Linux, Asianux	/etc/rc.d/rc5.d/S95VRTSralus.init /etc/rc.d/rc3.d/S95VRTSralus.init /etc/rc.d/rc2.d/S95VRTSralus.init /etc/rc.d/init.d/VRTSralus.init For example: rm /etc/rc.d/rc5.d/S95VRTSralus.init
Novell Open Enterprise Server 1.0/ SUSE Linux Enterprise Server 9 (32-bit only)	/etc/init.d/rc5.d/SxxVRTSralus.init /etc/init.d/rc3.d/SxxVRTSralus.init /etc/init.d/rc2.d/SxxVRTSralus.init /etc/init.d/VRTSralus.init For example: rm /etc/init.d/rc5.d/SxxVRTSralus.init
Novell Open Enterprise Server 2.0/ SUSE Linux Enterprise Server 10 (32-bit and 64-bit)	/etc/init.d/VRTSralus.init,start=2,3,5 /etc/init.d/VRTSralus.init For example: rm /etc/init.d/VRTSralus.init
Solaris	/etc/init.d/VRTSralus.init /etc/rc2.d/S91VRTSralus.init For example: rm /etc/rc2.d/S91VRTSralus.init
AIX	/etc/rc.d/VRTSralus.init /etc/rc.d/rc2.d/S95VRTSralus.init /etc/rc.d/rc3.d/S95VRTSralus.init /etc/rc.d/rc5.d/S95VRTSralus.init For example: rm /etc/rc.d/rc2.d/S95VRTSralus.init

See [“Manually uninstalling the Agent for Linux and Unix ”](#) on page 1446.

Starting the Agent for Linux and Unix daemon

If necessary, you can start the Agent for Linux and Unix (Linux and Unix Agent) daemon after the operating system starts.

See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.

To start the Agent for Linux and Unix daemon

- 1 Use a terminal session to connect to the Linux and Unix server as the root user.

- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

- 3 Start the Linux and Unix Agent daemon.

For example:

```
/etc/init.d/VRTSralus.init start
```

For Solaris:

```
/etc/init.d/VRTSralus.init start
```

For AIX:

```
/etc/rc.d/VRTSralus.init start
```

Stopping the Agent for Linux and Unix daemon

You can stop the Agent for Linux and Unix (Linux and Unix Agent) daemon.

See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

To stop the Agent for Linux and Unix daemon

- 1 Use a terminal session to connect to the Linux and Unix server as the root user.

- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

3 Stop the Linux and Unix Agent daemon:

For example:

```
/etc/init.d/VRTSralus.init stop
```

For Solaris:

```
/etc/init.d/VRTSralus.init stop
```

For AIX:

```
/etc/rc.d/VRTSralus.init stop
```

4 Restart the daemon when necessary.

Troubleshooting the Agent for Linux and Unix

If you experience problems with the Agent for Linux and Unix (Linux and Unix Agent) review the following questions and answers.

See “About the Agent for Linux and Unix” on page 1422.

Table O-7 Troubleshooting the Linux and Unix Agent

Question	Answer
Some characters do not appear correctly in the terminal session during the installation. What should I do?	This error occurs when the system location uses a non-English language character-set on the computer on which you install the Linux and Unix Agent. You can switch to another location setting of the same language to try to resolve this issue.
The Linux and Unix Agent installer is unable to install the Linux and Unix Agent. The following error is reported in the installralus log file. What should I do? VxIF::Error:: Unable to compress files. Hash(0x8711e8)->({GUNZIP} not found on <hostname>	To support the uncompressing of the Linux and Unix Agent platform-specific packages, you can install the GNU data compression utility. Install this utility on the computer on which you want to install the Linux and Unix Agent. The utility is available at the following URL: http://www.gzip.org

Table O-7 Troubleshooting the Linux and Unix Agent (continued)

Question	Answer
The Agent for Linux and Unix is installed on a Linux and Unix server in an NIS domain. Backup Exec is unable to browse resources on the server. What should I do?	Verify if the group line and the password line in the <code>nsswitch.conf</code> file are set to compatibility mode. If they are, then you must configure the <code>/etc/passwd</code> and <code>/etc/group</code> files. Refer to the <code>nsswitch.conf</code> man pages for additional information on how to configure the <code>nsswitch.conf</code> to use compatibility mode. Alternatively, change the password line and the group line to NIS files so that the Linux and Unix server validates the user through NIS. If the NIS server is unavailable or if the user is not found, the local files are used for validation.
I cannot load the Linux and Unix Agent. When I attempt to load the Linux and Unix Agent in console mode, <code>/beremote --log-console</code> shows the following message: ACE_SV_Semaphore_Complex: no space left on device. What should I do?	This issue occurs when the computer reaches its maximum limit on allowable semaphores. It can occur after an unexpected termination of the Linux and Unix Agent. When the Linux and Unix Agent unexpectedly terminates, it is unable to clean up some of the semaphore resources that it used. Other processes may have caused the use of semaphores to reach the limit. You must restart the computer to safely recover it from this condition. If other processes are running, it may not be feasible to restart the computer. Instead, you can use the commands that let you list and then remove all semaphores that are in use by the operating system. Be careful when you select semaphores to remove. Semaphores that are in use by the Linux and Unix Agent cannot be identified. If you remove semaphores of other programs that are in use, those programs can become unstable. To list semaphores, you can type the following command: <pre>ipcs -a</pre> To remove semaphores for each identifier that is listed, you can type the following command: <pre>ipcrm -s <id></pre>

Table O-7 Troubleshooting the Linux and Unix Agent (continued)

Question	Answer
I cannot load the Linux and Unix Agent. When I attempt to load the Linux and Unix Agent in console mode, /beremote --log-console shows the following message: Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory. What should I do?	This error indicates that the libstdc++.so.5 library is not in the /usr/lib directory. This library is necessary to let the Linux and Unix Agent start and function. To resolve this issue, install the libstdc++5 package. You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access: <pre>apt-get install libstdc++5</pre> For SUSE Linux Enterprise Server 11, run the following command: <pre>zypper install libstdc++5</pre>
On an Asianux operating system, Backup Exec displays an error that the Exec logon account for backup and restore options is not a part of the beoper group. What should I do?	Sometimes the /etc/group file is not correctly updated when a user is added to the beoper group. This leads to a failure of the getgrnam() Linux API. As a workaround, you should manually edit the /etc/group file and add the specific POSIX user name for the logon account in question to the beoper group. For more information on editing the /etc/group file, see the Asianux operating system documentation.

Table O-7 Troubleshooting the Linux and Unix Agent (continued)

Question	Answer
On any Linux and Unix computer on which the GNOME Virtual File System (GVFS) is installed and mounted, the GVFS cannot be browsed, backed up, or restored to.	On computers on which GVFS is installed, the .gvfs mount point is created for every user who logs on using the GNOME user interface. The mount point is created in the user's home directory. For example, if the user's logon name is John, then the directory appears as the following:
What should I do?	Echo <code>\$>ls -la /home/John</code> <pre>dr-x----- 2 John John 0 2009-06-16 18:16 .gvfs</pre> The output of the mount command appears as the following: <pre>"gvfs-fuse-daemon on /home/John/.gvfs type fuse.gvfs-fuse-daemon (rw,nosuid,nodev,user=John)".</pre> This mount point is created when the user logs on using the GNOME graphical user interface. The mount point is deleted when the user logs off. Note: : Logging on by using SSH or telnet does not show the mount point. A defect has been reported for the GVFS that even the superuser (root) does not have access to the file system. Refer to the following URL for more information: http://bugzilla.gnome.org/show_bug.cgi?id=560658 This defect prevents the Agent for Linux and Unix from running on the GVFS. Therefore, the GVFS cannot be browsed, backed up, or restored to. Other file systems on Linux and Unix computers on which GVFS is installed are also inaccessible. Additionally, on Linux and Unix computers on which GVFS is installed and mounted on the home folder of a particular user, browse operations, and backup and restore operations on or to the following file systems may fail: ■ SMBFS (commonly known as Samba) ■ Common Internet File System (CIFS) ■ Network File System (NFS) ■ ReiserFS As a resolution, try dismounting the GVFS, and retry the operation.

Table O-7 Troubleshooting the Linux and Unix Agent *(continued)*

Question	Answer
The Agent for Linux and Unix installer does not install the Perl Switch.pm module on 64 bit Ubuntu 14.04. What should I do?	Using the following steps, you must manually install Switch.pm before you attempt to install the Backup Exec Agent for Windows. ■ Open a terminal. ■ Type <code>cpan</code>. ■ Type <code>install Switch</code>. ■ Type <code>exit</code>. For Ubuntu 14.04, you must enable manual logins and add a root user using the following steps: ■ From <code>/usr/share/lightdm/lightdm.conf.d</code>, edit <code>50-ubuntu.conf</code>. ■ Add the following line: <code>greeter-show-manual-login=true</code> ■ Reboot the computer and add the root user on the login screen.

Backup Exec Remote Media Agent for Linux

This appendix includes the following topics:

- [About the Remote Media Agent for Linux](#)
- [How the Remote Media Agent for Linux works](#)
- [Requirements for the Remote Media Agent for Linux](#)
- [About open files and the Remote Media Agent for Linux](#)
- [About installing the Remote Media Agent for Linux](#)
- [Uninstalling the Remote Media Agent for Linux](#)
- [Starting the Remote Media Agent for Linux daemon](#)
- [Stopping the Remote Media Agent for Linux daemon](#)
- [About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers](#)
- [Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish](#)
- [Finding simulated tape library files](#)
- [About the Backup Exec operators \(beoper\) group for the Remote Media Agent for Linux](#)
- [About adding a Linux server as a Remote Media Agent for Linux](#)
- [Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux](#)

- [About creating storage device pools for devices attached to the Remote Media Agent for Linux](#)
- [Editing properties for the Remote Media Agent for Linux](#)
- [Deleting a Remote Media Agent for Linux from the Backup Exec list of servers](#)
- [Sharing a Remote Media Agent for Linux between multiple Backup Exec servers](#)
- [About backing up data by using the Remote Media Agent for Linux](#)
- [About restoring data by using the Remote Media Agent for Linux](#)
- [About the Tape Library Simulator Utility](#)
- [Creating a simulated tape library](#)
- [Viewing simulated tape libraries properties](#)
- [Deleting a simulated tape library](#)
- [Managing simulated tape libraries from the command line](#)
- [Command line switches for the Tape Library Simulator Utility](#)
- [Troubleshooting the Remote Media Agent for Linux](#)

About the Remote Media Agent for Linux

The Remote Media Agent for Linux lets you back up data from remote computers to the following devices:

- The storage devices that are directly attached to a Linux server.
- A simulated tape library on a Linux server.

You can add a Linux server to a Backup Exec server as a Remote Media Agent for Linux. You can then back up data from the Linux server or from supported remote computers to the devices that are attached to the Linux server. You can also create a virtual device on a Linux server on which the Remote Media Agent for Linux is installed. This virtual device emulates a SCSI tape library.

You can back up data from remote computers that have the following agents installed:

- Agent for Windows
- Agent for Linux
- Agent for Oracle on Linux or Windows Servers

The Remote Media Agent for Linux does not support backup for the following database agents and application agents:

- Agent for Microsoft SQL Server
- Agent for Microsoft Exchange Server
- Agent for Microsoft SharePoint
- Agent for Microsoft Active Directory
- Agent for Microsoft Hyper-V
- Agent for VMware
- Agent for Enterprise Vault

Note: In addition, the following items are not supported:

Backups that are enabled for Granular Recovery Technology (GRT).

Network-attached storage NDMP servers.

For information about the best practices to use Backup Exec Remote Media Agent for Linux (RMAL), refer to *Backup Exec Best Practices*.

See [“How the Remote Media Agent for Linux works”](#) on page 1457.

See [“About the Tape Library Simulator Utility”](#) on page 1475.

How the Remote Media Agent for Linux works

From the Backup Exec server, you can add a Linux server as a Remote Media Agent for Linux. You can then create backup, restore, and utility jobs on the Backup Exec server for the Linux server's storage devices.

When you back up data from a Windows server on which the Backup Exec Agent for Windows is installed to a storage device that is attached to a Remote Media Agent for Linux, the data connection is established directly between the two systems. In this case, the Backup Exec server does not play any role in data transfer.

If the storage device that is attached to the Remote Media Agent for Linux is an OpenStorage device, a direct connection is required between the data and the device to enable client-side deduplication to occur. If a direct connection is not available, then the data flows through the Backup Exec server to perform server-side deduplication.

If you use Backup Exec Central Admin Server feature, you can share a Remote Media Agent for Linux computer between multiple Backup Exec servers. Sharing

can be enabled when you add the Remote Media Agent for Linux. You can select new Backup Exec servers to share the Remote Media Agent for Linux or remove the sharing ability from the Backup Exec servers at any time.

See [“Sharing storage devices”](#) on page 557.

Job performance increases because data travels from the remote computers to the devices that are attached to the Linux server. This increase is especially apparent if the Backup Exec server is located at a different site than the Remote Media Agent for Linux computer and the remote computers.

The Remote Media Agent for Linux does not have a user interface. You use the administration console on the Backup Exec server to manage the jobs and devices on the Remote Media Agent for Linux. The Backup Exec server maintains job logs, catalogs, job histories, alerts, and notifications.

See [“Requirements for the Remote Media Agent for Linux”](#) on page 1458.

See [“About installing the Remote Media Agent for Linux”](#) on page 1459.

See [“About adding a Linux server as a Remote Media Agent for Linux”](#) on page 1468.

See [“About the Tape Library Simulator Utility”](#) on page 1475.

Requirements for the Remote Media Agent for Linux

The following items are required to install the Remote Media Agent for Linux:

- You must have superuser privileges on the Linux server.
- The Linux server must have Perl 5.8.8 or later installed.

You can find a list of compatible devices, operating systems, platforms, and applications in the Backup Exec Hardware and Software Compatibility Lists.

It is recommended that you use the Secure Shell (SSH) protocol when you push-install the Remote Media Agent for Linux to remote servers. You must enable SSH before you install the Remote Media Agent for Linux.

Note: Some versions of Linux may require that you install the libstdc++.so.5 package.

See [“About installing the Remote Media Agent for Linux”](#) on page 1459.

About open files and the Remote Media Agent for Linux

The Remote Media Agent for Linux uses advanced open file and image technologies that are designed to alleviate issues that are sometimes encountered during backup operations, such as backing up open files.

After you make file and folder selections and the job is submitted for backup, the Remote Media Agent for Linux automatically makes a snapshot of the volume or volumes. Making a snapshot of a volume provides a point-in-time record of the data. When the Remote Media Agent for Linux creates a snapshot, it uses snapshot technologies to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. During the backup, files can be open and data can be changed.

Remote Media Agent for Linux supports Simple, Logical Volume Manager (LVM), and RAID volume configurations.

See [“Requirements for the Remote Media Agent for Linux”](#) on page 1458.

About installing the Remote Media Agent for Linux

Use the Backup Exec installation media to do the following:

- Install Remote Media Agent for Linux on a local Linux server.
- Push-install the Remote Media Agent for Linux to one or more remote Linux servers.

If you push-install the Remote Media Agent for Linux, the RSH (Remote Shell) is used by default. It is recommended that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install the Remote Media Agent for Linux. Refer to your operating system documentation for more information about SSH.

Before you install the Remote Media Agent for Linux, you should review the requirements:

See [“Requirements for the Remote Media Agent for Linux”](#) on page 1458.

When you install the Remote Media Agent for Linux, Backup Exec creates the beoper group and adds root as a member. Any Linux user that you add to the beoper group gets the permissions necessary to back up and restore the Linux servers.

However, if Backup Exec detects an NIS server during the Remote Media Agent for Linux installation, then the beoper group is not created. You must create the beoper group manually on the Linux servers.

After the installation completes, you must add the Linux server as a Remote Media Agent on the Backup Exec server. Then, you can send jobs to the devices that are attached to the Linux server.

See [“Installing the Remote Media Agent for Linux”](#) on page 1460.

See [“Creating the Backup Exec operators \(beoper\) group manually for the Remote Media Agent for Linux”](#) on page 1468.

See [“About adding a Linux server as a Remote Media Agent for Linux”](#) on page 1468.

See [“About the Backup Exec operators \(beoper\) group for the Remote Media Agent for Linux”](#) on page 1467.

Installing the Remote Media Agent for Linux

You can install Remote Media Agent for Linux on a local Linux server or push-install it to one or more remote Linux servers.

See [“About installing the Remote Media Agent for Linux”](#) on page 1459.

Note: You must unzip the RALUS_RMALS_<version number>.gz file on a Linux server. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Remote Media Agent for Linux

- 1 At a Linux server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install the Remote Media Agent for Linux.
- 3 Navigate to the following path on the installation media:
<Linux>
- 4 Copy the RALUS_RMALS_<version number>.gz file in this directory to a directory on the local server.
- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_<version number>.gz
```

- 6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_<version number>.tar
```


7 Start the **installrmal** script.

For example:

```
./installrmal
```

8 Do one of the following:

To install on a local server	Press Enter .
To install to one remote server	Type the name, IP address, or fully qualified domain name of a Linux server.
To install to multiple remote servers	Type the names, IP addresses, or fully qualified domain names of the Linux servers. Leave a space between each identifier.

9 After the installer checks for a valid Linux operating system during the initial system check, press **Enter**.

10 Review the package installation summary, and then press **Enter**.

11 After the system installation requirements check completes, press **Enter**.

12 Start the prerequisites check by pressing **Enter**.

13 Type the name, IP address, or fully qualified domain name of the Backup Exec server (directory host) that you want to use this Remote Media Agent.

14 Type any names, IP addresses, or fully qualified domain names of additional Backup Exec servers that you want to use this Remote Media Agent.

15 Do one of the following:

If the server name, IP address, or fully qualified domain name is correct	Press Enter to continue the installation.
If you want to change a server name, IP address, or fully qualified domain name	Type N , press Enter , and then change the information.

16 Start the NIS server scan by pressing **Enter**.

17 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected	The Remote Media Agent for Linux installer cannot create the beoper group. You must create it manually after the Remote Media Agent for Linux installation is complete. Continue with the next step.
If an NIS server is not detected	Use the installer to create the beoper group. Do the following in the order listed: ■ To let the installer create the beoper group, type y. ■ To select the next available Group ID, type n. ■ To add the root user account to the beoper group, type y. ■ Continue with the next step.

- 18 Start the installation by pressing **Enter**.
- 19 After installation completes, press **Enter** to start the post-installation configurations and installation of SymSnap drivers.
- 20 Press **Y** to automatically start the Beremote service; otherwise, press **N** to start the service later.
- 21 After the configuration process completes, press **Enter** to save the installation log to the following file:
`/var/tmp/vxif/installrmalsummary file number/installrmal.log`
- 22 If the RMAL installer did not create a beoper group, you must create it.
See [“Creating the Backup Exec operators \(beoper\) group manually for the Remote Media Agent for Linux”](#) on page 1468.
- 23 Start the Agent for Linux daemon.
See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.
- 24 Add the Linux server as a Remote Media Agent.
See [“About adding a Linux server as a Remote Media Agent for Linux ”](#) on page 1468.

Uninstalling the Remote Media Agent for Linux

You must have the Backup Exec installation media to uninstall the Remote Media Agent for Linux.

Before you uninstall the Remote Media Agent for Linux, you should note the location of the simulated tape library files. Then, you can delete all of the simulated tape library files after the uninstall operation completes. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Finding simulated tape library files”](#) on page 1466.

To uninstall the Remote Media Agent for Linux

- 1 On the Linux server, place the Backup Exec installation media in the appropriate device.
- 2 Log on as root to the server from which you want to uninstall the Remote Media Agent for Linux.

- 3 Navigate to the following path on the installation media:

<Linux>

- 4 Start the **uninstallrmal** script.

For example:

```
./uninstallrmal
```

- 5 Do one of the following:

To uninstall the Remote Media Agent for Linux from one server	Type the name, IP address, or fully qualified domain name of the Linux server.

To uninstall the Remote Media Agent for Linux from multiple servers	Type the names, IP addresses, or the fully qualified domain names of the Linux servers. Leave a space between each identifier.

- 6 Press **Enter**.
- 7 After the Remote Media Agent for Linux package check completes successfully, press **Enter**.

- 8 When you are prompted to uninstall the Remote Media Agent for Linux packages, press **Enter** to save the uninstall summary and log on to the following location:

`/var/tmp/vxif/uninstallrmalsummary file number.log`

- 9 Manually delete the simulated tape library files.

Starting the Remote Media Agent for Linux daemon

If necessary, you can start the Remote Media Agent for Linux daemon after the operating system starts.

See [“Stopping the Remote Media Agent for Linux daemon”](#) on page 1464.

To start the Remote Media Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

`/etc/init.d/`

For example:

`cd /etc/init.d/`

- 3 Start the Remote Media Agent for Linux daemon.

For example:

`/etc/init.d/VRTSralus.init start`

Stopping the Remote Media Agent for Linux daemon

You can stop the Remote Media Agent for Linux daemon.

See [“Starting the Remote Media Agent for Linux daemon”](#) on page 1464.

To stop the Remote Media Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

`/etc/init.d/`

For example:

`cd /etc/init.d/`

- 3 Stop the Remote Media Agent for Linux daemon.

For example:

```
/etc/init.d/VRTSralus.init stop
```

- 4 Restart the Remote Media Agent for Linux daemon when necessary.

About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers

When you connect to a Remote Media Agent for Linux computer from the Backup Exec server, you must establish trust between the Backup Exec server and the Remote Media Agent for Linux computer. You must also establish trust if you want to configure a remote Remote Media Agent for Linux computer to perform client-side deduplication.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 953.

See [“Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers”](#) on page 1465.

Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers

You can add one or more Remote Media Agent for Linux computers to the Backup Exec list of servers. When you add the Remote Media Agent for Linux, you must establish a trust between the Backup Exec server and the remote Linux computers to ensure secure communication.

To establish trust and add a Remote Media Agent for Linux computer to the Backup Exec list of servers

- 1 On the **Storage** tab, click **Configure Storage**, and then select **Network storage**.
- 2 Click **Next**.
- 3 Select **Backup Exec Remote Media Agent for Linux**, and then click **Next**.
- 4 Follow the on-screen prompts.

See [“About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers”](#) on page 1465.

Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish

You can specify additional Backup Exec servers to which the Remote Media Agent for Linux can publish information.

Each Backup Exec server to which the Remote Media Agent for Linux publishes information appears in the Backup Exec list of servers.

To add additional Backup Exec servers to which the Remote Media Agent for Linux can publish information

- 1 Use a text editor to open the following file:
 /etc/VRTSralus/ralus.cfg
- 2 Add the following string:
 Software\Veritas\Backup Exec For Windows\Backup
 Exec\Engine\Agents\Agent Directory List *unique identifier number* = *IP
 address or DNS name of Backup Exec server*
- 3 Save and close the file.
- 4 Move to the Backup Exec server to which the Remote Media Agent for Linux is publishing itself and add the Remote Media Agent for Linux server to the **Servers** list.

See [“About the list of servers on the Backup and Restore tab”](#) on page 150.

Finding simulated tape library files

Before you uninstall the Remote Media Agent for Linux, you should note the location of the simulated tape library files. Then, after you uninstall the Remote Media Agent for Linux, you can delete all of the simulated tape library files. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Uninstalling the Remote Media Agent for Linux ”](#) on page 1463.

See [“About the Tape Library Simulator Utility”](#) on page 1475.

To find simulated tape library files

- 1 Log on as root to the server on which you want to find the simulated tape library files.
- 2 Navigate to the following directory that contains the Tape Library Simulator:
/opt/VRTSralus/bin

For example:

```
cd /opt/VRTSRalus/bin
```

- 3 Start the **mktls** utility to list the simulated tape library files and folders.
For example:

```
/opt/VRTSralus/bin/mktls -l
```

- 4 Write down the locations of the directories for the simulated tape library files.

About the Backup Exec operators (beoper) group for the Remote Media Agent for Linux

The Backup Exec operators (beoper) group contains the names of the users who have permission to back up and restore the Linux servers.

When you install the Remote Media Agent for Linux, Backup Exec creates the (beoper) group and adds root as a member. Any Linux user that you add to the beoper group gets the necessary permission to back up and restore the Linux servers.

However, if an NIS server is detected during the Remote Media Agent for Linux installation, Backup Exec cannot create the (beoper) group. You must create the (beoper) group manually on the Linux servers on which you want to install the Remote Media Agent for Linux. You must create the (beoper) group before you start backup and restore operations. Otherwise, connections fail between the Linux servers and the Backup Exec server.

Before the members of the (beoper) group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators \(beoper\) group manually for the Remote Media Agent for Linux”](#) on page 1468.

See [“Backup Exec logon accounts”](#) on page 759.

Creating the Backup Exec operators (beoper) group manually for the Remote Media Agent for Linux

If the Remote Media Agent for Linux installation detects an NIS server during installation, you must create a beoper group on each Linux server on which you install the Remote Media Agent for Linux.

See [“About the Backup Exec operators \(beoper\) group for the Remote Media Agent for Linux”](#) on page 1467.

Note: Ensure that you understand how to set security for groups on Linux servers before you assign a Group ID for the beoper group.

Table P-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux server on which you want to install the Remote Media Agent for Linux. If the Linux server is in an NIS domain, navigate to the NIS domain's group file.	See the NIS documentation for information on how to add a group to a NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux server.	See the operating system documentation for more information about how to add users to a group
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See “Creating a Backup Exec logon account” on page 762.

About adding a Linux server as a Remote Media Agent for Linux

After you add the Linux server as a Remote Media Agent for Linux, and if you purchased and installed the Central Admin Server feature, you can share the

storage devices that are connected to the Remote Media Agent for Linux computer with other Backup Exec servers.

See [“Adding a Linux server as a Remote Media Agent for Linux”](#) on page 1469.

See [“Sharing storage devices”](#) on page 557.

See [“About the Central Admin Server feature”](#) on page 1321.

Adding a Linux server as a Remote Media Agent for Linux

Use the following steps to add a Linux server as a Remote Media Agent for Linux.

See [“About adding a Linux server as a Remote Media Agent for Linux ”](#) on page 1468.

To add a Linux server as a Remote Media Agent for Linux

- 1 On the **Storage** tab, click **Configure Storage**.
- 2 Select **Network storage**, and then click **Next**.
- 3 Select **Backup Exec Remote Media Agent for Linux**, and then click **Next**.
- 4 Follow the on-screen prompts to trust the Remote Media Agent for Linux and to restart the services.

See [“Sharing storage devices”](#) on page 557.

Remote Media Agent for Linux options

You must provide information when you add a Linux server as a Remote Media Agent for Linux to a Backup Exec server.

See [“About adding a Linux server as a Remote Media Agent for Linux ”](#) on page 1468.

Table P-2 Add Remote Media Agent for Linux options

Item	Description
Server name	Specifies the name of the Linux server that you want to add as a Remote Media Agent for Linux. If the Backup Exec Central Admin Server feature is installed in your environment, use the host name or fully qualified domain name of the Linux server. That is, use the name of the Linux computer that appears when you browse for backup selections. If you use the IP address, Backup Exec cannot distinguish which device path to use for jobs.

Table P-2 Add Remote Media Agent for Linux options (*continued*)

Item	Description
Port number	Lists the port to use for communications between the Backup Exec server and the Remote Media Agent for Linux. If you change the port number, you must edit the services file in the /etc directory on the Linux server, and update the NDMP entry. See “Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux” on page 1471. Ensure that this port is open in any firewalls that exist between the Remote Media Agent for Linux and the Backup Exec server. Use a port number that is not in use by another application or service. The default port is 10000.
Description	Displays a description that you choose.
Logon Account	Indicates the logon account for the Remote Media Agent for Linux. The default logon account is the system logon account for the Backup Exec server.
Have Backup Exec use ICMP ping operations to detect the server	Lets the Backup Exec server use ICMP ping operations to locate the Linux server. You can turn off this option in environments where ping requests are blocked. This option is enabled by default.
Logon account	Indicates the Backup Exec logon account that you want to use to log on this server. See “Backup Exec logon accounts” on page 759.

See [“About creating storage device pools for devices attached to the Remote Media Agent for Linux ”](#) on page 1471.

See [“About backing up data by using the Remote Media Agent for Linux”](#) on page 1474.

Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux

You can change the port that Backup Exec uses to communicate with the Remote Media Agent for Linux.

To change the port for communications between the Backup Exec server and the Remote Media Agent for Linux

- 1 On the computer on which the Remote Media Agent for Linux is installed, use a text editor to open the services file in the /etc directory.

For example:

```
vi /etc/services
```

- 2 Search the file for an entry that is similar to the following:

ndmp 10000/tcp

- 3 Do one of the following:

If this entry exists

Change the port number to the port number that you want to use.

If this entry does not exist

Do the following in the order listed:

- At the end of the file, type `ndmp`, and then press **Tab**.
- Type the port number that you want NDMP to use, and then type `/tcp`.
- Press **Enter**.

- 4 Save the file, and then exit the editor.

- 5 Restart the Agent for Linux daemon.

See [“Starting the Remote Media Agent for Linux daemon”](#) on page 1464.

About creating storage device pools for devices attached to the Remote Media Agent for Linux

Remote Media Agents may reside in different physical locations. To reduce network traffic and increase job performance, you can create separate storage device pools for Remote Media Agents that are located at different sites.

See [“Creating storage device pools”](#) on page 527.

See [“About the Tape Library Simulator Utility”](#) on page 1475.

Editing properties for the Remote Media Agent for Linux

You can edit the properties for a Remote Media Agent for Linux.

To edit properties for the Remote Media Agent for Linux

- 1 On the **Storage** tab, right-click a Remote Media Agent for Linux server.
- 2 Click **Details**.

See [“Remote Media Agent for Linux properties”](#) on page 1472.

Remote Media Agent for Linux properties

You can view or edit properties for a Remote Media Agent for Linux server.

See [“Editing properties for the Remote Media Agent for Linux ”](#) on page 1472.

The following table lists Remote Meida Agent for Linux server properties:

Table P-3 Remote Media Agent for Linux properties

Item	Description
Name	Displays the name, IP address, or fully qualified domain name of the Remote Media Agent for Linux.
Port	Displays the port that is used for communications between the Backup Exec server and the Remote Media Agent for Linux.
Backup Exec server status	Displays the status of the Backup Exec server. Backup Exec server status includes Online, Pause, Unavailable, and Offline.
Description	Displays a description of the Remote Media Agent for Linux. You can edit this description.

Table P-3 Remote Media Agent for Linux properties (continued)

Item	Description
Enable ICMP ping operations for Backup Exec to detect the Remote Media Agent	Lets Backup Exec communicate with the Remote Media Agent for Linux. You can turn off this option in environments where ping requests are blocked. This option is enabled by default.
Host ID	Displays the identifier number that the Remote Media Agent for Linux generates.
System version	Displays the version of the operating system that runs on the Remote Media Agent for Linux.
Logon account	Indicates the logon account for the Remote Media Agent for Linux. Click Change to select or create another logon account.

Deleting a Remote Media Agent for Linux from the Backup Exec list of servers

Use the following steps to delete a Remote Media Agent for Linux from the Backup Exec list of servers.

To delete a Remote Media Agent for Linux from a Backup Exec

- 1 On the **Storage** tab, right-click a Remote Media Agent for Linux.
- 2 Click **Delete**.
- 3 Click **Yes**.

See [“Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers”](#) on page 1465.

Sharing a Remote Media Agent for Linux between multiple Backup Exec servers

If the Central Admin Server feature is installed, you can select Backup Exec servers to share a Remote Media Agent for Linux server. When you add a Remote Media Agent for Linux server, the Backup Exec server that you used to add the device is automatically selected for sharing.

See [“Sharing storage devices”](#) on page 557.

To share a Remote Media Agent for Linux between multiple Backup Exec servers

- 1 On the **Storage** tab, under **All Storage**, right-click a Remote Media Agent for Linux server that you want the Backup Exec servers to access.
- 2 Select **Share**.
- 3 Under **Server**, select the Backup Exec servers that you want to use with the Remote Media Agent for Linux.
- 4 Click **OK**.
- 5 Restart the Backup Exec services on the Backup Exec servers that you selected in step 3.

See [“Starting and stopping Backup Exec services”](#) on page 771.

About backing up data by using the Remote Media Agent for Linux

Create a backup job for the Remote Media Agent for Linux from the Backup Exec server.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 205.

See [“Adding a stage to a backup definition”](#) on page 220.

About restoring data by using the Remote Media Agent for Linux

Create a restore job for the Remote Media Agent for Linux from the Backup Exec server.

Note: Use devices that are attached to the Backup Exec server to restore data from the tapes that other applications created. Remote Media Agent for Linux supports only Microsoft Tape Format (MTF) media.

See [“Methods for restoring data in Backup Exec”](#) on page 235.

About the Tape Library Simulator Utility

The Tape Library Simulator Utility lets you create a virtual device on a hard disk or on any mounted volume on a Linux server. This virtual device emulates a SCSI tape library. The Remote Media Agent for Linux must be installed on the server.

When you run the Tape Library Simulator Utility, you are prompted for the following information:

- The number of slots that you want to allocate to this library.
- The location or path for the library.

The Tape Library Simulator Utility then creates the media for the simulated tape library. To ensure that each media has a unique name, the Tape Library Simulator Utility creates a bar code label for each media. You cannot rename these bar code labels. However, you can add a unique media description.

The simulated tape library emulates an Advanced Intelligent Tape (AIT) media type. This media type is seldom used, so it helps you distinguish between a physical robotic library and a simulated tape library. The simulated media also has an AIT media type label.

The format of the files that are written to the simulated tape library is similar to the file format of backup-to-disk files. However, you cannot copy or move files between simulated tape libraries and disk-based storage.

You can add the simulated tape library to Backup Exec storage device pools.

See [“Creating storage device pools”](#) on page 527.

To use the Tape Library Simulator Utility, you must have a minimum of 500 MB of available space on the Linux server. The available space includes hard disk space, flash drives, and USB drives. If there is not enough space, the jobs fail with an end-of-media error. You must either create available disk space or you must direct the jobs to another volume, and then start the jobs again.

A simulated tape library does not support all of the tasks that are available for physical robotic libraries.

Note: You cannot run fast cataloging of the tapes that have become inseparable. Backup Exec automatically reverts to slow cataloging. The following message is written in the catalog job's job log:

"An error was encountered while attempting to read the Media Based Catalog's Set Map data. Attempting to catalog the media by reading the data area of each backup set."

This problem occurs because an incomplete on-tape catalog is written to the tape at the end of a backup job. As a result, subsequent backup jobs cannot find the tape's set map, and flag the tape as unreadable.

See ["About storage operation jobs"](#) on page 534.

See ["Creating a simulated tape library"](#) on page 1476.

Creating a simulated tape library

Create a simulated tape library on a server on which the Remote Media Agent for Linux is installed. You must create the simulated tape library on a hard disk or on a mounted volume.

See ["About the Tape Library Simulator Utility"](#) on page 1475.

To create a simulated tape library

- 1 At the Remote Media Agent for Linux computer, stop the Agent for Linux daemon.

See ["Stopping the Agent for Linux and Unix daemon"](#) on page 1449.

- 2 Navigate to the following path that contains the Tape Library Simulator Utility:

```
</opt/VRTSralus/bin>
```

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Start the **mktls** utility.

For example:

```
./mktls
```

- 4 Select **Create a new simulated tape library**, and then press **Enter**.

- 5 Enter the appropriate information.

See ["Simulated Tape Library options"](#) on page 1477.

- 6 Exit the utility.
- 7 Restart the Agent for Linux daemon.
 See “[Starting the Agent for Linux and Unix daemon](#)” on page 1449.
- 8 On the Backup Exec server, restart the Backup Exec services.
 See “[Starting and stopping Backup Exec services](#)” on page 771.

Simulated Tape Library options

When you create a simulated tape library, you must provide a directory path and the number of slots for the library.

See “[Creating a simulated tape library](#)” on page 1476.

Table P-4 Simulated Tape Library options

Item	Description
Directory Path	Type the path of the directory for the simulated tape library. You can enter up to 512 characters. If the path does not exist, the Tape Library Simulator Utility creates it.
Number of Slots	Select the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.

See “[Viewing simulated tape libraries properties](#)” on page 1477.

Viewing simulated tape libraries properties

You can use the Veritas Tape Library Simulator Utility to view information about a simulated tape library and its contents.

To view simulated tape library properties

- 1 On the Remote Media Agent for Linux computer, stop the Agent for Linux daemon.

See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.

- 2 Navigate to the following directory that contains the Tape Library Simulator Utility:

`/opt/VRTSralus/bin`

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Start the **mktls** utility.

For example:

```
./mktls
```

- 4 Select **View an existing simulated tape library**.
- 5 Move your cursor to the simulated tape library that you want to view, and then press **Enter**.
- 6 Press **Enter** again to view the simulated tape library properties.
- 7 Type **Q** to exit the utility.
- 8 Restart the Agent for Linux daemon.

See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

Simulated tape library properties

You can view the properties of a simulated tape library.

Table P-5 Simulated tape library properties

Item	Description
Number of drives	Displays the number of drives for this simulated tape library. A simulated tape library can have only one drive. This drive is not configurable.

Table P-5 Simulated tape library properties (*continued*)

Item	Description
Number of slots	Displays the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.
Tape capacity	Displays the tape capacity. The default capacity is 100 gigabytes.
Directory path	Displays the directory path where the simulated tape library exists.

Deleting a simulated tape library

You can use the Tape Library Simulator Utility to delete a simulated tape library. You must then manually delete the content of the simulated tape library files, and then delete the directories that contain these files.

To delete a simulated tape library

- 1 At the Remote Media Agent for Linux computer, stop the Agent for Linux daemon.

See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.
- 2 Navigate to the following directory that contains the Tape Library Simulator:

/opt/VRTSralus/bin/

For example:


```
cd /opt/VRTSralus/bin/
```
- 3 Start the **mktls** utility:

For example:


```
./mktls
```
- 4 Select **View an existing simulated tape library**.
- 5 Select the simulated tape library that you want to delete.
- 6 When you are prompted, delete the simulated tape library.
- 7 Exit the utility.
- 8 Restart the Agent for Linux daemon.

See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

- 9 Find the simulated tape library files, and then manually delete them.
See [“About the Tape Library Simulator Utility”](#) on page 1475.
- 10 On the Backup Exec server, restart the Backup Exec services when it is convenient.
See [“Starting and stopping Backup Exec services”](#) on page 771.

Managing simulated tape libraries from the command line

You can use the command line to create a simulated tape library. Create a simulated tape library on a hard disk or on any mounted volume on the Remote Media Agent for Linux computer. From the command line, you can also view and delete simulated tape libraries.

To manage simulated tape libraries from the command line

- 1 At the Remote Media Agent for Linux computer, stop the Agent for Linux daemon.
See [“Stopping the Agent for Linux and Unix daemon”](#) on page 1449.
- 2 Navigate to the following directory that contains the Tape Library Simulator Utility:

`/opt/VRTSralus/bin`

For example:

`cd /opt/VRTSralus/bin`
- 3 Start the **mktls** utility with the appropriate parameter switches.
See [“Command line switches for the Tape Library Simulator Utility”](#) on page 1480.
- 4 Start the Agent for Linux daemon.
See [“Starting the Agent for Linux and Unix daemon”](#) on page 1449.

Command line switches for the Tape Library Simulator Utility

You can use command line switches to manage simulated tape libraries. For example, the following command line creates a simulated tape library with 10 slots that is located at `/TLS2/Testing`.

```
./mktls -s10 -p/TLS2/Testing
```

See [“Managing simulated tape libraries from the command line”](#) on page 1480.

Table P-6 Command line switches for the Tape Library Simulator Utility

Switch	Description
-p<path>	Specifies the path to the directory for the simulated tape library. If the path does not exist, the utility creates it. The maximum path size is 512 characters.
-s<number of slots>	Specifies the number of slots for this simulated tape library. The number of slots can range from one to 50. The default number is 20.
-r	Prevents the information from displaying.
-l	Lists the simulated tape libraries that exist for the Remote Media Agent for Linux.
-d -p<path>	Specifies the path of the simulated tape library that you want to delete.
-h	Displays the online Help.

Troubleshooting the Remote Media Agent for Linux

If there are issues with Remote Media Agent for Linux, review the following questions and answers.

Table P-7 Troubleshooting the RMAL

Question	Answer
Remote Media Agent for Linux does not detect my attached device. What should I do?	First, ensure that Backup Exec and the Remote Media Agent for Linux support the device. You can find a list of compatible devices in the Backup Exec Hardware Compatibility List. If the device is listed on the hardware compatibility list, ensure the following: ■ The operating system detects the device ■ The device is listed in /proc/scsi/scsi If the operating system can detect the device, ensure that the device is listed in /etc/VRTSralus/TILDBG.TXT.
My Backup Exec server does not display the devices that are attached to my Remote Media Agent. What should I do?	Try the following procedures: ■ Ensure that the Agent for Linux daemon is running. If it is not running, start the daemon, and check that power for the server is enabled, and that all cables are properly attached. ■ Ensure that the Remote Media Agent for Linux properties are set to the correct port, and that ICMP ping operations are enabled. ■ Ensure that the Backup Exec services are restarted after a Remote Media Agent is added to the Backup Exec server. The available devices should be displayed under the Remote Media Agent for the Linux node. See “Editing properties for the Remote Media Agent for Linux ” on page 1472. See “Starting the Agent for Linux and Unix daemon” on page 1449.
Why don't my remote devices appear in any of the storage device pools that are created by Backup Exec?	By default, Backup Exec does not include remote devices in the storage device pools it creates. It is recommended that you create a separate storage device pool for the devices that are attached to each Remote Media Agent. See “About creating storage device pools for devices attached to the Remote Media Agent for Linux ” on page 1471.

Table P-7 Troubleshooting the RMAL (*continued*)

Question	Answer
Remote Media Agent for Linux won't run on the remote computer. What should I do?	Ensure that the Remote Media Agent for Linux is installed on a supported version of Linux. You can find a list of compatible operating systems, platforms, and applications in the Backup Exec Software Compatibility List. If you install the Remote Media Agent for Linux to an unsupported version of Linux, the Remote Media Agent for Linux is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Agent for Linux component. This component is installed with the Remote Media Agent for Linux. To use the Agent for Linux component to back up the Linux server, do the following: ■ Edit the <code>ralus.cfg</code> file. ■ In the string <code>Software\Veritas\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0</code>, change 0 to 1. See "Editing configuration options for Linux and Unix computers" on page 1432. - See "Running the begather utility to troubleshoot Backup Exec components on Linux servers" on page 875.
I cannot load the Remote Media Agent for Linux. When I attempt to load the Remote Media Agent for Linux in console mode, <code>/beremote --log-console</code> shows the following message: Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory. What should I do?	This error indicates that the <code>libstdc++.so.5</code> library is not in the <code>/usr/lib</code> directory. This library is necessary to let the Remote Media Agent for Linux start and function. To resolve this issue, install the <code>libstdc++.so.5</code> package. You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access: <pre>apt-get install libstdc++5</pre> For SUSE Linux Enterprise Server 11, run the following command: <pre>zypper install libstdc++5</pre>

Glossary

ADAMM (Advanced Device and Media Management)	A Backup Exec database that automates the tracking of tape cartridge media. ADAMM expires the backup sets that are stored on tape cartridge media according to the associated media set.
administration console	The user interface that allows you to run Backup Exec operations. The user interface can be run from the Backup Exec server or from a remote computer.
agent	A component that allows computers such as Microsoft SQL Server to interact with the Backup Exec server.
Agent for Windows	A Backup Exec system service that runs on Microsoft Windows computers and allows remote backup and restore of those computers.
alert	An event in Backup Exec that usually requires some form of user interaction or acknowledgment.
alert category	A group of one or more events that occur in Backup Exec and that can generate an alert. Examples of alert categories include Job Success, Install Warning, and Database Maintenance Failure.
alert source	A source that can generate an alert. Alert sources include jobs, media, storage devices, and computers.
alert type	The classification of an alert that lets you determine the severity of the alert. Alert types include Error, Warning, Information, and Attention Required.
allocated media	The tape cartridge media that are associated with a media set and that have current append and overwrite protection periods.
append period	The length of time that data can be added to tape cartridge media. The append period starts when the first backup job is written to the media.
audit log	A running history of all actions that are performed in Backup Exec. An entry into the log is created each time an action occurs that is configured to display in the audit log.
backup definition	A container for any backup selections, job templates, and stages that you specify. Backup Exec combines the job templates with the backup selections to create backup jobs. If you specify a stage, then that additional task is also run.
Backup Exec server	The computer on which Backup Exec is installed and where the Backup Exec services are running.

Backup Exec server pool	A feature of the Backup Exec Central Admin Server Option that lets you group managed Backup Exec servers in a pool to which you can restrict backup jobs.
Backup Exec service account	A user account that is configured for the Backup Exec system services. It contains a user name and password and provides the rights to log on as a service and to act as a Backup Exec administrator.
backup method	An option that you select when you run a backup job to specify a full, differential, or incremental backup.
backup set	A collection of data that is backed up from a single source of content. For example, a single source of content can be a server or it can be a Microsoft Exchange dataset. If you select multiple sources of contents, Backup Exec creates multiple backup sets.
backup strategy	The procedures that you implement for backing up your network. Backup strategies include what methods of backup are performed and when backups are performed.
baseline	The first backup job to run in a synthetic backup. The baseline backup runs one time only and backs up all of the files on the selected computer. A full backup is assembled, or synthesized, from a baseline backup and the subsequent incremental backups.
catalog	A database that Backup Exec creates during a backup operation. When you select data to restore, Backup Exec uses the catalog information to find the restore selections and the storage devices on which they reside.
central administration server	A Backup Exec server on which the Central Admin Server Option (CASO) is installed. In a CASO environment, the central administration server provides centralized administration, delegated job processing, and load balancing functionality for managed Backup Exec servers.
centralized catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are kept on the central administration server.
cloud storage	An online storage location on multiple virtual servers to which you can back up data.
common encryption key	A type of encryption key that anyone can use to back up data using encryption and to restore encrypted data.
custom error-handling rule	An error-handling rule that you can define for a specific error code in an error category. When a job fails with the error code that is associated with the custom error-handling rule, the retry options and the final job disposition are applied to the job.
data discovery	A Backup Exec feature that allows the detection of new backup content within a Windows domain.

data lifecycle management (DLM)	An automated disk reclamation process that Backup Exec uses to delete expired backup sets that are on disk-based storage. The disk space is then free for use by new backup sets. DLM deletes backup sets from disk-based storage after the amount of time expires that you specified when you created the backup job. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-component of a server, even if the backup sets are expired.
differential	A backup method that includes all files that have changed since the last full backup.
disk storage	A location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data.
distributed catalog	A catalog location in the Central Admin Server Option. Image files in the catalog are distributed to the central administration server from every managed Backup Exec server. These distributed files are small because they do not contain the entire catalog. They contain only information about the backup set. The history files, which contain detailed information about the backup set, remain on the managed Backup Exec server.
error-handling rule	A default or custom rule that sets retry options and the final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you place the job on hold until you can fix the error.
event	An action that occurs during a Backup Exec operation, such as a job cancellation.
full	A backup method that includes all of the files that you select for backup.
granular restore	A restore of individual items from a backup for which you enable the Granular Recovery Technology option.
GRT (Granular Recovery Technology)	A backup option that is available with some Backup Exec agents. Granular Recovery Technology lets you restore individual items from database backups. A separate backup of the individual items is not required for you to recover one item.
imported media	The media that are created by a product other than this installation of Backup Exec, but are in storage devices in the Backup Exec environment.
incremental	A backup method that backs up only the files that have changed since the last full or incremental backup.
job delegation	A process by which jobs are distributed by a central administration server to available storage devices on managed Backup Exec servers. Job delegation is only available with the Central Admin Server Option.
job	An operation that has been scheduled for processing by the Backup Exec server. Jobs contain source or destination information, settings, and a schedule. Types

	of jobs include backup, restore, data discovery, reports, test run, and storage operations.
job history	A list of completed and failed backup, restore, and storage operation jobs.
job log	A log that contains the results of a job. It is created when the job runs. You can review the job log for job errors and job details.
job template	A collection of settings that Backup Exec uses to create jobs. For example, backup job settings can include encryption, scheduling options, or notifications. When a backup job is run, Backup Exec combines the job template with the backup selections to create a backup job.
legacy backup-to-disk folder	A storage device used in versions prior to Backup Exec 2012 that you could create to back up data to a folder on a hard disk. For later versions, these legacy backup-to-disk folders are read-only. It is recommended that you use disk storage devices instead.
load balancing	A feature in Backup Exec that automatically distributes jobs among any available storage devices in a storage device pool. Also a feature of the Backup Exec Central Admin Server Option in which jobs are automatically distributed from a central administration server to multiple managed Backup Exec servers for processing among the various storage devices.
logon account	An account that stores the credentials of a Windows user account and that enables Backup Exec to manage user names and passwords. It can be used to browse data sources or to process jobs.
managed Backup Exec server	A Backup Exec server that is managed by a central administration server. Managed Backup Exec servers are responsible for the actual processing of backup and restore jobs in a Central Admin Server Option environment. Managed Backup Exec servers are only available with the Backup Exec Central Admin Server Option.
media ID	A unique internal label that Backup Exec assigns to each media used in Backup Exec. The ID keeps statistics for each media. The media ID cannot be erased or changed.
media label	A label used to identify media. Backup Exec can assign the label automatically, but you can rename it. If the media was first used in a library with a barcode reader, the media label will already have a barcode label.
media overwrite protection level	A global setting in Backup Exec that lets you specify whether to overwrite scratch, imported, or allocated tape cartridge media regardless of the media's overwrite protection period.
media rotation	A strategy that determines when tape cartridge media can be reused, or rotated back into use, by Backup Exec. Common examples of a media rotation strategy are Son, Father/Son, and Grandfather/Father/Son.

media set	A set of rules that apply to tape cartridge media that are associated with a media set. These rules specify append periods, overwrite protection periods, and vaulting periods.
media vault	A logical representation of the actual physical location of tape cartridge media, such as a special media room, a scratch bin, or an offsite location.
mixed backup	A backup definition that contains more than one backup method for multiple data types.
offhost backup	A feature of the Backup Exec Advanced Disk-based Backup Option that enables the backup operation to be processed on a Backup Exec server instead of on the remote computer, or host computer. Moving the backup from the remote computer to a Backup Exec server enables better backup performance and frees the remote computer as well.
Offline Tape Cartridge Media vault	A location on the Storage tab that displays the tape cartridge media that are on-site but are not in tape drives, robotic libraries, or media vaults. Media are automatically moved to the offline vault if you use Backup Exec to remove media from a tape drive or robotic library.
Online Tape Cartridge Media vault	A location on the Storage tab that displays the tape cartridge media that are available in tape drives or robotic libraries. You cannot add or move media to the online media vault. Backup Exec does that automatically.
overwrite protection period	The length of time that data is retained on a specific tape cartridge media before being overwritten (unless the media is erased, formatted, moved to scratch media, or if the media overwrite protection level is set to None). The overwrite protection period is measured from the last time data was appended to the media.
preferred server configuration	A collection of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where data is replicated between multiple servers.
recyclable media	Tape cartridge media that are assigned to a media set but have expired data overwrite protection periods.
remote administrator	The Backup Exec user interface (Administration Console) that is run on remote computers.
replicated catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are replicated from the managed Backup Exec server to the central administration server.
restricted encryption key	A type of encryption key that anyone can use to back up data using encryption. Only the key owner or a user with knowledge of the pass phrase can restore data that was encrypted with a restricted encryption key.
retired media	Tape cartridge media that has been taken out of service, usually because of an excessive number of errors. Media that is retired is available for restore jobs but

	not for backup jobs. Media must be retired before it can be deleted. If you want to use media that has been deleted, Backup Exec recognizes it as imported media. You must catalog retired media before you can restore from it.
scratch media	Tape cartridge media that are not associated with a media set and that can be overwritten. Scratch media includes new or blank media, erased media, and media moved from another group.
simulated tape library	A tape library that emulates an Advanced Intelligent Tape (AIT) media type and has the AIT media type label. A simulated tape library is created by the Tape Library Simulator.
stage	An additional task that you can run with a backup job, such as duplicating a copy of the backup data to disk storage.
storage device	A disk storage device, disk cartridge, robotic library drive, stand-alone drive, virtual drive, removable storage drive, cloud-based storage device, or other type of data storage that is supported by Backup Exec.
storage device pool	A group of similar types of storage devices that enables load-balancing of Backup Exec jobs.
synthetic backup	A feature of the Advanced Disk-based Backup Option that enables a full backup to be assembled, or synthesized, from a baseline and subsequent incremental backups.
Tape Library Simulator	A utility that lets you create a virtual device on a hard disk or on any mounted volume on a computer on which the Backup Exec Remote Media Agent for Linux is installed. The virtual device that is created is called a simulated tape library.
true image restore	A feature of the Advanced Disk-based Backup Option that enables Backup Exec to restore the contents of directories to what they were at the time of any full or incremental backup. Restore selections are made from a view of the directories as they existed at the time of the particular backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not unnecessarily restored and then overwritten.
UMI (Unique Message Identifier)	A unique code that is associated with an error reported in the job log, or on some alerts. These codes contain hyperlinks that you can click to go to the Technical Support Web site. You can access technical notes and troubleshooting tips that are related to a specific error.
virtual disk	A logical disk that you configure on a storage array to provide storage to the Backup Exec server.

Index

A

- about
 - AWS FSx 1416
 - Azure Files 1416
 - Backup Exec Cloud Deduplication 388
 - forever incremental backups 940
- active alerts
 - responding to 310
- Active Directory
 - backing up in Exchange 1178
- Active File Exclusion 1182
- active jobs
 - about managing and monitoring 261
 - canceling 265
 - holding 266
 - removing a hold 267
 - statuses 289
 - viewing job activity 265
 - viewing properties 265
- Add
 - Remote Media Agent for Linux 1468
- adding tapes by importing media 568
- administration console
 - overview 113
 - role in backup process 35
- Advanced Disk-based Backup feature
 - about 1382
 - baseline
 - setting 1383
 - best practices for offhost backup 1390, 1393
 - host computer
 - defined 1388
 - offhost backup options 1386, 1392
 - offhost backup overview 1388
 - setting default options 1385
 - transportable snapshots
 - defined 1388
 - true image restore
 - overview 1386
- Advanced Open File
 - configuring options for backup jobs 669
- Agent for Hyper-V
 - about removing an instantly recovered virtual machine 1107
 - adding a Hyper-V host 1070
 - and GRT 1087
 - backing up 1074
 - Full catalogs 1091
 - highly available virtual machines 1090
 - how byte count is calculated for full catalog 1092
 - installation overview 1068
 - installing Agent for Windows 1072
 - Instant GRT 1091
 - notes 1068
 - overview 1064
 - protecting Microsoft application data 1087
 - push-installing Agent for Windows 1073
 - requirements 1066
 - restoring 1093
 - setting default backup options 1083
 - viewing details about resources 1071
 - with GRT and vhd/x files 1068
- Agent for Linux
 - about establishing trust 953
 - push-installing 1423
 - using SSH 1423
- Agent for Linux and Unix
 - about backing up 1439
 - about exclusions from backup 1431
 - backup job options 1439
 - beoper group, defined 1427
 - configuration options in the ralus.cfg file 1432
 - configuring the ralus.cfg file 1430
 - creating the beoper group 1427
 - default options 1443
 - editing configuration options in the ralus.cfg file 1432
 - editing default options 1443
 - establishing trust relationship 1428
 - installing 1424
 - publishing to Backup Exec servers 1429
 - requirements 1423

- Agent for Linux and Unix *(continued)*
 - restore options 1442
 - restoring 1441
 - runtime scripts 1447
 - saving the installation log 1424
 - starting the Linux and Unix Agent daemon 1449
 - stopping the Linux and Unix Agent daemon 1449
 - troubleshooting 1450
 - uninstalling 1445
 - uninstalling manually 1446
- Agent for Microsoft Active Directory
 - about 1311
 - about restoring individual objects 1317
 - Granular Recovery Technology (GRT)
 - overview 1313
 - passwords 1318
 - recreating purged objects 1317
 - requirements 1312
 - tombstones 1317
- Agent for Microsoft SharePoint
 - about 1204
 - adding a farm 1207
 - backing up SharePoint data 1207
 - disabling or enabling communication between web servers and Backup Exec 1216
 - disaster recovery for SharePoint 1218
 - installing 1205
 - overview 1204
 - requirements 1205
 - restoring SharePoint data 1214
 - system requirements 1205
 - using with SharePoint Server and SharePoint Foundation 1206
- Agent for VMware
 - adding VMware vCenter and ESX/ESXi hosts 1008
 - backing up 1013
 - backing up Microsoft application data 1028
 - backup defaults 1022
 - dynamic inclusion 1027
 - Granular Recovery Technology
 - about 1028
 - GRT requirements 1028
 - installing 1008
 - installing Agent for Windows 1010
 - instant recovery overview 1044
 - overview 1005
 - requirements 1005
- Agent for VMware *(continued)*
 - restoring resources 1034
 - selecting transport method for VMDK file 1020, 1026
 - unsupported characters 1005
 - unsupported GRT configurations 1028
- Agent for Windows
 - about 951
 - about establishing trust 953
 - Backup Exec Agent Utility 954
 - hardware requirements 952
 - installation methods 68
 - installing in an Active Directory network 75
 - installing on Hyper-V virtual machines 1072
 - installing on virtual machines 1010
 - installing updates 74
 - installing using a command script 83
 - installing using the command prompt 79
 - licenses 952
 - publish to Backup Exec servers 958
 - push-installing to remote computers 68
 - push-installing to VMware virtual machines 1011
 - push-installing to Hyper-V virtual machines 1073
 - requirements 952
 - stopping and starting 953
 - uninstalling using a command script 84
 - uninstalling using command prompt 82
- agent-based backup 139
- agents
 - upgrading 106
- Alert History by Backup Exec server report 808
- Alert History report 808
- alerts
 - categories 303
 - clearing informational alerts 311
 - configuring categories 319
 - configuring defaults 323
 - configuring groups for notification 316
 - copying text 309
 - defined 303
 - deleting from alert history 308
 - deleting recipients 318
 - disabling pop-ups 324
 - enabling email and text messages 314
 - enabling pop-ups 324
 - filters 309
 - notification 312

- alerts *(continued)*
 - responding to 310
 - sending job complete notification 321
 - setting up notification 311
 - severity 303
 - showing on the Home tab 306
 - SNMP traps 325
 - stopping notification for recipient 322
 - viewing job log 310
 - where to find 304
- All Servers server group
 - about 153
- allocated media
 - overwriting 506
- Alternate location
 - setting an SDR 889
- Amazon cloud
 - requirements 369
- Amazon cloud storage
 - configure 369
- append period
 - changing 503
 - defined 495
 - editing 501
- archive bit
 - using to determine backed up status 200
- Arctera Knowledge Base
 - searching 868
- Arctera Update
 - about 101
 - running manually 102
 - scheduling automatic updates 101
- audit log
 - about 774
 - configuring 774
 - removing entries 775
 - saving to a file 776
 - viewing 775
- Audit Log report 810
- automatic exclusion of SQL data during volume level
 - backups 1127
- AWS FSx and Azure Files
 - add server 1418
 - back up 1418
 - best practices 1418
 - notes 1417
 - recommendation 1417
 - restore 1419
- AWS FSx and Azure Files
 - pre-requisites 1417
- Azure file share
 - limitation 1420
- B**
 - back up and delete the files method
 - freeing disk space 202
 - backp definition
 - creating 158
 - Backup
 - using the Remote Media Agent for Linux 1474
 - backup
 - creating 158
 - overview 158
 - VMware virtual machine 1013
 - Backup and Restore tab
 - list of servers 150
 - backup definition
 - creating 161
 - creating from an existing backup definition 161
 - defined 158
 - editing 205
 - excluding selections 179
 - including selections 182
 - one-time 158
 - selecting data 171
 - Backup Exec
 - Lock Console 116
 - locking and unlocking the console 116
 - overview
 - how it works 35
 - Backup Exec Agent Utility
 - activity status
 - viewing 956
 - command line applet 963
 - switches 964
 - using 963
 - database access
 - configuring 960
 - Event Viewer
 - open 954
 - job template name for DBA-initiated jobs 1237
 - Linux
 - configure Oracle instance on 1239
 - port
 - configure for Oracle operations 1240
 - publish to Backup Exec servers 958

- Backup Exec Agent Utility *(continued)*
 - publishing
 - adding Backup Exec servers 958
 - editing Backup Exec server information 959
 - removing Backup Exec servers 960
 - Real Application Cluster (RAC)
 - publish to Backup Exec server 1235
 - refresh interval
 - setting 957
 - Registry Editor
 - open 954
 - security
 - removing certificate 963
 - Services
 - open 954
 - starting 955
 - starting automatically 957
 - update credentials for Linux instances 1236
 - view status 955
 - Windows
 - configure Oracle instance on 1233
- Backup Exec Cloud Deduplication
 - command-line tool 394
 - configure 398
 - delete 400
 - notes 389
 - requirements 390
 - WORM 393
- Backup Exec diagnostic application
 - generating 871
- Backup Exec Migrator
 - about 1287
 - about retrieving Enterprise Vault data 1305
 - about staged migrations 1292
 - about the Backup Exec Backup Sets view 1304
 - Backup Exec server
 - working with 1299
 - best practices 1308
 - communicating with Enterprise Vault 1302
 - configuring 1297
 - data migration process 1292
 - Enterprise Vault retention periods 1296
 - events
 - about 1293
 - how it works 1288
 - log file location 1294
 - logs
 - about 1293
- Backup Exec Migrator *(continued)*
 - migrated files
 - about deleting 1296
 - Migrator for Enterprise Vault options 1301
 - requirements 1288
 - retrieving Enterprise Vault data 1306
 - troubleshooting 1309
- Backup Exec server 35
 - viewing properties 777
- Backup Exec services
 - changing service account credentials 771
 - changing startup options 773
 - stopping and starting 771
- Backup Exec settings
 - changing preferences 698
 - database maintenance and security 701
 - DBA-initiated jobs 746
 - discover data to back up 713
 - Granular Recovery Technology (GRT)
 - options 745
 - network and security 718
- backup job
 - creating 158, 161
 - creating from an existing backup definition 161
 - deduplication 996
 - editing 205
 - excluding selections 179
 - excluding selections globally 179
 - including selections 182
 - one-time 158
 - pre/post commands 674
 - preparing for 138
 - required user rights 150
 - running the next scheduled instance 204
 - selecting data 171
- Backup Job Success Rate report 811
- backup methods
 - about 189
 - advantages and disadvantages 192
 - configuring 197
 - delete selected files and folders after successful
 - backup 202
 - differential 191
 - duplicate 190
 - full 190
 - incremental 192
 - specific types of data 194
- backup network
 - overview 716

- backup or restore job
 - debugging enabled 273, 278
- Backup Recommendations report 811
- Backup Resource Success Rate report 812
- backup selections
 - about 171
 - changing the order 176
 - critical system components 175
 - multiple servers or applications 174
 - using fully qualified domain names 176
- backup sets
 - about 359
 - about duplicating 222
 - about verifying 227
 - automatically deleting expired 352
 - cataloging 260
 - changing the expiration dates 362
 - duplicating 223
 - expiring 362
 - prevent from expiring 363
 - releasing from retention 364
 - retaining 363
 - verifying 228
 - viewing contents 365
 - viewing properties 365
- Backup Sets by Media Set report 812
- Backup Size By Resource report 813
- backup strategies
 - increase throughput with Agent for Windows 951
- barcode labels
 - default 511
 - overview 511
 - robotic library support 511
- barcode media, finding in drives and portals 561
- baseline
 - setting for synthetic backup 1383
- beoper group
 - Agent for Linux and Unix, about 1427
 - creating 1427
- Boot managers
 - restoring in SDR 916
- buffer count
 - setting for tape drives 476
- buffer size
 - setting for disk cartridge devices 348
 - setting for tape drives 475
- byte count
 - how it is calculated for catalog operations 1092

- byte count *(continued)*
 - incorrect 859

C

- calendar
 - excluding dates 695
 - viewing scheduled backup jobs 220
- CAS
 - alerts 1358
 - Backup Exec server
 - changing to a managed Backup Exec server 1338
 - Backup Exec server pool
 - adding managed Backup Exec servers 1365
 - creating 1364
 - deleting 1365
 - overview 1363
 - removing a managed Backup Exec server 1366
 - selecting for backup 1364
 - Backup Exec Utility
 - running 1380
 - catalog locations 1344
 - central administration server
 - setting for a managed Backup Exec server 1338
 - centralized restore
 - multiple storage devices 1367
 - overview 1366
 - changing to a central administration server 1337
 - communication thresholds 1357
 - deleting managed Backup Exec server 1339
 - disabling communications 1358
 - disaster recovery 1377
 - enabling communications 1358
 - installing 1326
 - installing across a firewall 1333
 - installing managed Backup Exec server 1327
 - job delegation 1360
 - managed Backup Exec server
 - configuration options 1331
 - viewing properties 1372
 - network interface cards
 - using any available 1360
 - network traffic
 - reducing 1343
 - notifications 1358
 - overview 1321

CAS (continued)

- pausing managed Backup Exec server 1370
 - recovering failed jobs 1369
 - requirements 1323
 - restoring data from multiple devices 1367
 - resuming paused managed Backup Exec server 1370
 - settings for managed Backup Exec servers 1346
 - starting Backup Exec services 1371
 - status 1357
 - stopping Backup Exec services 1371
 - storage and media data 1324
 - troubleshooting 1378
 - uninstalling Backup Exec from central administration server 1380
 - uninstalling Backup Exec from managed Backup Exec server 1381
 - upgrading 1335
- catalog
- configuring full catalog options 662
 - defined 254
 - editing default options 255
 - full catalog 254
 - Instant GRT 254
 - levels 258
 - media with encrypted backup sets 520
- catalog operation errors
- DLT tape drive hangs 853
- cataloging
- media 560
- catalogs
- using to determine backed up status 200
- centralized catalogs in CAS 1344
- centralized restore
- best practices 1367
- CHECKCATALOG utility 1124
- CHECKDB utility 1124
- checkpoint restart
- about 672
 - configuring 672
 - configuring default settings 673
 - considerations 673
 - supported technologies 673
- cleaning a drive 567
- cleaning slots
- defining for robotic libraries 485
- client-side deduplication
- overview 995

Cloud deduplication storage device

- configuring 390
 - deleting 390
 - requirements 390
- Cloud Storage Summary report 813
- CloudConnect Optimizer
- about 408
 - configure job 409
 - delete job 411
 - edit job 410
- command line
- installing Backup Exec 90
 - installing Remote Administrator 86
 - switches for installation 91
- command prompt
- uninstalling Agent for Windows 82
- common encryption keys 732
- compact view 119
- completed jobs
- statuses 289
- configuration settings
- copying to another server 776
- Configure Storage Wizard
- overview 538
- configuring
- holidays 695, 698
- consistency check options
- Exchange Agent 1184, 1191
 - SQL Agent 1123
- continuing Exchange backup if consistency check fails 1184, 1191
- control connection with the Remote Media Agent for Linux 1457
- conversion to virtual machines
- adding a stage for 613
 - after a backup job 605
 - backup methods 589
 - conversion of disks larger than 2TB 589
 - from point in time 620
 - how backup selections are processed 589
 - one-time conversion 622
 - options 589
 - overview 589
 - requirements 597
 - setting default options 628
 - simultaneously with backup job 599
 - Windows Server 2012 Hyper-V host 589
- Copy text 117

- credentials
 - creating for backup sources 214
 - creating for jobs 214
 - replacing for a backup source 213
 - replacing for jobs 213
 - testing for backup sources 212
 - testing for jobs 212
- Credentials pane
 - about 212, 215
 - deleting backup sources 215
- critical system components
 - about 184
 - restore scenarios 187
 - selecting 175
- D**
- Daily Device Utilization report 814
- damaged tape media
 - removing 517
- data connection to remote computers 1457
- data lifecycle management, overview 352
- database maintenance and security
 - about 701
 - configuring 701
- Database snapshots
 - SQL 1138, 1146
- DBA-initiated jobs
 - creating a template 747
 - deleting a template 748
 - duplicate job settings 755
 - editing a template 748
 - general options 754
 - network options 754
 - storage options 749
 - templates 746
- Debug Monitor 875
- debugging enabled
 - backup or restore job 273, 278
- deduplication
 - Windows 170
- Deduplication Disk and Cloud Deduplication Device
 - Summary report 815
- deduplication disk storage
 - changing logon account password 990
 - overview 982
- deduplication disk storage devices
 - disaster recovery 1001
 - editing properties 985
 - requirements 972
- deduplication disk storage devices (continued)
 - restoring 1001
- Deduplication feature
 - about backing up 996
 - about copying deduplicated data to tapes 1000
 - changing logon account password for
 - deduplication disk storage 990
 - client-side deduplication overview 995
 - copying data between OpenStorage devices or
 - deduplication disk storage devices 997
 - deduplication disk storage overview 982
 - deduplication disk storage properties 985
 - deduplication methods for agents 971
 - direct access
 - editing properties 992
 - selecting storage devices 991
 - disaster recovery of deduplication disk
 - storage 1001
 - disaster recovery of OpenStorage devices 1002
 - installing 980
 - OpenStorage device overview 419
 - OpenStorage device properties 422
 - overview 969
 - requirements 972
 - restoring a deduplication disk storage
 - device 1001
 - sharing devices 995
 - with encryption 1000
- Deduplication Summary report 816
- default backup options
 - Hyper-V 1083
- Default options
 - Simplified Disaster Recovery
 - settings 890
- default options
 - Agent for VMware 1022
 - conversion to virtual machines 628
 - NDMP feature 1412
- default preferred configuration settings for tape
 - drives 476
- deleting
 - tape media 518
- destination Backup Exec server
 - adding 776
 - importing a list 776
- Details pane
 - Hyper-V 1071
- Device Summary report 817

- devices
 - adding iSCSI-attached 472
 - OpenStorage overview 419
 - reconnecting USB tape devices 472
- diagnostic file
 - command line switches 872
- differential backups
 - about 191
 - advantages and disadvantages 193
- direct access
 - editing properties 992
 - selecting storage devices 991
- DirectCopy to tape
 - copying data 230
 - overview 229
- directories
 - about including and excluding for NDMP servers 1406
- disaster preparation
 - Disaster Preparation Plan (DPP) 131
 - Exchange Server 1202
 - hardware protection 132
 - off-site storage 132
 - overview 131
- Disaster recovery
 - alternate path in SDR 891
 - setting path locations
 - disaster recovery information file 890
- disaster recovery
 - deduplication disk storage 1001
 - Exchange Server 1202
 - Microsoft SharePoint 1218
 - OpenStorage devices 1002
- discover data to back up
 - about 713
 - adding servers 713
 - configuring 713
- disk cartridge storage
 - about 345
- disk storage
 - creating 334
 - editing properties 338
- Disk Storage Summary report 817
- disk-based storage
 - about 330
- distributed catalogs in CAS 1344
- DLT tape
 - drive hangs when cataloging 853

- domain controller
 - installing by redirected restore 249
 - restoring System State 248
- duplication between OpenStorage devices or deduplication disk storage devices 997
- dynamic inclusion
 - for Hyper-V 1086

E

- editions of Backup Exec
 - listed and described 35
- eject media 567
 - after job completes 657
- email notification
 - configuring 312
- encrypted files
 - about cataloging media 520
- encrypted SQL database restore 1146
- encryption
 - about 729
 - hardware 731
 - restoring encrypted SQL databases 1146
 - software 730
 - types 729
 - with deduplication 1000
- encryption keys
 - 128-bit AES 729
 - 256-bit AES 729
 - common 732
 - creating 734
 - deleting 738
 - encryption types 729
 - managing 733
 - overview 731
 - pass phrases 732
 - replacing 735
 - restoring encrypted data 253
 - restricted 732
- Enterprise Vault Agent
 - about redirecting a restore job 1283
 - about restoring 1277
 - About restoring individual files and folders 1282
 - automatic redirection of Enterprise Vault components 1282
 - available backup methods 1272
 - Backup Exec Migrator
 - about 1287
 - about deleting migrated files 1296

- Enterprise Vault Agent *(continued)*
 - Backup Exec Migrator *(continued)*
 - about events 1293
 - about logs 1293
 - about retrieving Enterprise Vault data 1305
 - about staged migrations 1292
 - about the Backup Exec Backup Sets view 1304
 - best practices 1308
 - communicating with Enterprise Vault 1302
 - configuring 1297
 - data migration process 1292
 - Enterprise Vault retention periods 1296
 - how it works 1288
 - log file location 1294
 - Migrator for Enterprise Vault options 1301
 - requirements 1288
 - retrieving Enterprise Vault data 1306
 - troubleshooting 1309
 - VxBSA logs 1293
 - working with a Backup Exec server 1299
 - Backup Exec server
 - log file location 1294
 - logs 1293
 - best practices 1287
 - collections
 - configuring 1298
 - vault store partition properties 1298
 - installing 1270
 - migration
 - vault store partition properties 1303
 - non-operational state 1277
 - Partition Recovery Utility
 - about 1306
 - finding an archive ID 1307
 - log file location 1294
 - logs 1293
 - requirements 1306
 - running 1307
 - troubleshooting 1309
 - ready-to-use state 1277
 - redirection options 1284
 - requirements 1269
 - restore options 1279
 - selecting a backup method 1271
- Environment Check
 - results 42
 - running before installing or upgrading 42
- error codes
 - Unique Message Identifier viewing 282
- error-handling rules
 - creating 286
 - custom rules
 - defined 285
 - custom rules for recovered jobs 288
 - default rules
 - defined 285
 - deleting a custom rule 287
 - enabling for a failed job 287
 - enabling or disabling 287
 - overview 285
 - recovered jobs custom rule 286
- Error-Handling Rules report 818
- ESX/ESXi host, adding 1008
- Event Recipients report 819
- Exchange Agent
 - Active Directory
 - backing up 1178
 - automatic exclusion of files during volume level backups 1182
 - backing up
 - Exchange 1182
 - recommended selections 1178
 - backup methods 1187, 1189, 1194–1195
 - best practices 1174
 - Database Availability Group 1185
 - databases
 - configuring 1196
 - disaster recovery 1202
 - Exchange in a Database Availability Group 1192
 - Exchange Web Services
 - overview 1179
 - excluding files during volume level backups 1182
 - Granular Recovery Technology (GRT)
 - overview 1179
 - requirements for 1163
 - setting for backup 1190, 1196
 - installation 1170
 - Internet Information Service (IIS) metabase
 - backing up 1178
 - mailbox access requirements 1175
 - offhost backup
 - with Granular Recovery Technology (GRT) 1179

Exchange Agent *(continued)*
 overview 1161
 protecting Exchange using VSS 1180
 redirecting data 1196
 requirements 1161
 restore of individual items
 requirements 1163
 restore requirements 1196
 restoring data from snapshot backups 1196
 restoring data to server 1196
 restoring individual public folder messages from
 tape 1196
 services account 1162
 snapshot backup
 configuring 1180
 snapshot technology 1180
 strategies for backing up 1176
 system state
 backing up 1178
 troubleshooting snapshot and offhost jobs 1182
 volume level backups
 automatic exclusion of files 1182

Exchange Web Services
 using with the Exchange Agent 1179

exclude dates
 deleting dates 697
 exporting dates to another server 698
 importing a list of dates 695
 selecting dates for all backups 695
 selecting dates for individual backup jobs 218

exporting expired media 576
 exporting media 576

F

Failed Backup Jobs report 820

failed jobs
 retry 272

FAT
 partition 859

father/son media rotation strategy 523

file history
 enabling for NDMP feature 1412

files
 about including and excluding for NDMP
 servers 1406

files and folders options
 configuring for backup jobs 683

filters
 for alerts 309

firewall

Backup Exec ports 726
 browsing systems through 726
 enabling a SQL instance behind 729
 using Backup Exec with 725

forever incremental backup

backing up 947
 CAS-MBES 948
 catalog 946
 differences 141
 DLM 945
 GRT 946
 how it works 941
 limitations 950
 notes 948
 recommendations 949
 retain 945
 schedule 944
 supported storage 946

formatting media 566

full backups

about 190
 advantages and disadvantages 192

Full catalog

and Hyper-V virtual machine backups 1091

full catalog

how byte count is calculated for Hyper-V 1092

full catalog operation

configuring 662
 for GRT-enabled jobs 739

G

GDPR Guard

about 841
 BEMCLI commands 843
 best practices 848
 how to block 846
 restoring blocked items 847
 supported types of data 845
 troubleshooting 865

Google cloud

requirements 372

Google cloud storage

configure 373

grandfather media rotation strategy 523

Granular Recovery Technology

about using with VMware 1028
 requirements for VMware 1028
 unsupported configurations for VMware 1028

- Granular Recovery Technology (continued)
 - with the Agent for Hyper-V 1087
- Granular Recovery Technology (GRT)
 - about restoring individual items 739
 - Exchange data 1179
 - offhost backup 1179
 - recommended devices for 744
 - requirements 743
 - setting default options 745
 - staging 743
 - using Exchange Web Services 1179
- Group Policy Object, configuring 78
- groups
 - configuring to receive notifications 316
- GRT
 - about using with VMware 1028
 - requirements for VMware 1028
 - unsupported configurations for VMware 1028

H

- hardware
 - protection in case of disaster 132
 - troubleshooting 851
- high water count
 - setting for tape drives 476
- highly available virtual machines
 - about backing up and restoring 1090
- Home tab
 - about 120
 - configuring 120
 - Layout items 120
 - restoring the default configuration 120
 - Support items 120
 - System Health items 120
- Hyper-V
 - instant recovery 1098, 1104
- Hyper-V host
 - adding to list of servers 1070
- Hyper-V instant recovery
 - best practices 1109
 - limitations 1103
 - remove 1107
 - removing 1108
 - requirements 1103

I

- IBM computers
 - recovering with Simplified Disaster Recovery 915
- imported media
 - overwriting 506
- imported tape media
 - labeled by Backup Exec 511
- importing media 568
- In-Memory optimized tables
 - restoring SQL databases that contain In-Memory optimized tables 1146
- include dates
 - creating 217
- incremental backups
 - about 192
 - advantages and disadvantages 193
- initializing a robotic library 564
- installation
 - additional features 58
 - Agent for Windows 68
 - Agent for Windows from command prompt 79
 - Agent for Windows with command script 83
 - checklist 41
 - command line switches 91
 - configuring Group Policy Object 78
 - creating a transform 75
 - Environment Check
 - overview 42
 - from command line 90
 - Media Agent for Linux 1460
 - Microsoft SQL Server Express 44
 - Migration Report 109
 - NDMP feature 1400
 - overview 39
 - parameter files
 - creating 97
 - post-installation tasks 109
 - pre-upgrade checklist 108
 - push-installing Agent for Windows 68
 - push-installing to remote computers 60
 - Remote Administrator 84
 - Remote Administrator from command line 86
 - system requirements 45
 - typical 48
 - uninstalling Agent for Windows using command prompt 82
 - uninstalling Agent for Windows with command script 84

- installation *(continued)*
 - uninstalling Backup Exec 110
 - uninstalling features from local Backup Exec server 111
 - updates to Agent for Windows 74
 - Windows Management Instrumentation performance counter 328
 - Windows Management Instrumentation SNMP provider 329
 - installation log 98
 - Agent for Linux and Unix 1424
 - installation overview
 - Agent for Hyper-V 1068
 - installation parameter file
 - creating 97
 - Installation Summary Report 99
 - viewing 99
 - installed updates
 - viewing 103
 - installing
 - SharePoint Agent 1205
 - to an existing Microsoft SQL Server 2008 instance 44
 - Instant GRT
 - configuring 662
 - for GRT-enabled jobs 739
 - Instant Recovery
 - Hyper-V 1098
 - instant recovery
 - creating job for a VMware virtual machine 1049
 - removing a VMware virtual machine 1052
 - VMware notes 1048
 - VMware overview 1044
 - VMware requirements 1047
 - Internet Information Services (IIS) metabase
 - backing up 1178
 - inventory
 - robotic libraries when Backup Exec services start 480
 - IPv4 718
 - IPv6 718
 - iSCSI-attached devices
 - adding 472
- J**
- job activity 265
 - job defaults
 - backing up multiple servers or applications 700
 - job defaults *(continued)*
 - backup jobs 640
 - changing 640
 - exclude dates 695
 - exclude selections 179
 - schedule for rule-based and run now jobs 692
 - job history 263
 - about duplicating 222
 - about verifying 227
 - deleting jobs 275
 - duplicating 225
 - overview 274
 - retry only failed resources 277
 - running a job 276
 - verifying 228
 - viewing 275
 - job log 279
 - configuring default options 283
 - finding text 279
 - linking to technical support web site 282
 - printing 281
 - saving 281
 - viewing from an alert 310
 - with vertical applications 282
 - job monitor 261, 263
 - job progress indicators
 - displaying 699
 - job queue
 - holding 268
 - removing the hold 269
 - job status 289
 - job status and recovery 295
 - job template
 - defined 158
 - jobs
 - canceling 265
 - changing priority for scheduled 270
 - configuring error-handling rules 286
 - deleting from Job History 275
 - deleting scheduled 272
 - holding 266
 - holding the job queue 268
 - holiday scheduling 695, 698
 - managing and monitoring 261
 - removing a hold 267
 - removing hold on the job queue 269
 - retry only failed resources 272
 - running from Job History 276
 - running scheduled job 270

- jobs *(continued)*
 - sending notification when complete 321
 - setting status and recovery options 295
 - viewing the job log 279
- Jobs Summary report 821

L

- labeling media 509
- labeling tape media
 - imported media label 511
 - renaming 512
 - using barcode labels 511
- Least Free Space First storage device in pool 530
- Library Expansion feature
 - SCSI addresses for hardware 479
 - setting up hardware 479
- license contract information
 - about 104
 - managing customer numbers 105
- license information
 - viewing 104
- licenses 39
 - Agent for Windows 952
- licensing
 - instance-based usage 125
- list of servers
 - about 150
 - adding discovered servers 715
 - adding servers 151
 - removing servers 152
 - server groups 153
- list view 119
- local Backup Exec server
 - breaking connection with 88
- local server properties
 - about viewing 777
 - viewing 777
- logon accounts
 - about 759
 - changing default 768
 - changing the password 765
 - checking 712
 - copying to another server 769
 - creating 762
 - default
 - about 760
 - deleting 766
 - editing 763
 - replacing 766

- logon accounts *(continued)*
 - restricted 761
 - scheduling a test 712
 - system logon account 760
 - testing 770
- logon information
 - copying to another server 769

M

- mailbox access requirements for Exchange 1175
- Malware scan
 - best practices 239
 - configure 238
 - limitation 240
 - options 783
- Malware Scan Summary report 821
- managed Backup Exec server
 - changing settings 1346
 - copying jobs to 1361
 - installing 1327
 - network interface card
 - using any available 1360
 - pools 1363
 - upgrading 1335
- Managed Backup Exec Servers report 822
- master database (SQL)
 - backup 1123
- media
 - erasing 518
 - overwrite options 506
 - overwriting allocated or imported 506
 - retired
 - defined 493
 - scratch
 - defined 493
- Media Agent for Linux
 - beoper group 1460
 - installing 1460
- Media Audit report 824
- Media Errors report 825
- media ID
 - defined 509
- media overwrite protection level
 - defined 505
- Media Required for Recovery report 825
- media rotation
 - strategies
 - father/son 523
 - grandfather 523

- media rotation *(continued)*
 - strategies *(continued)*
 - son 523
- media set
 - creating 499
 - default 491
 - deleting 505
 - renaming 505
 - vault rule properties 515
- Media Summary report 826
- Media Vault Contents report 827
- media vaults, about 513
- messages
 - error 857
- Microsoft 365
 - about 429
 - backup tenant data 433
 - best practices 459
 - CAS-MBES 447
 - catalog operations 447
 - configure tenant 430
 - delete tenant 433
 - install 429
 - limitations 451
 - notes 448
 - parallel streams and job settings 679
 - recommendations 458
 - requirements 429
 - restore tenant data 439
 - selections 433
 - update tenant 432
 - User selection 434
 - view tenant 431
 - workflow 429
- Microsoft 365 tenant
 - viewing properties 778
- Microsoft SharePoint data
 - backing up 1207
 - restoring 1214
- Microsoft SQL Server
 - installing 44
- Microsoft Virtual Hard Disk files
 - about managing 170
- Migration Report 109
- modified time
 - using to determine backed up status 199
- Most Free Space First storage device in pool 530
- Move Media to Vaultreport 827

- MSDE
 - components
 - installed with Backup Exec 44
- Multi-stream backup
 - virtual machines 1033

N

- named transaction
 - restore up to 1146
- NDMP feature
 - adding NDMP servers 1400
 - backing up NDMP servers 1403
 - duplicate backed up data 1408
 - how to use patterns 1406
 - installing 1400
 - overview 1399
 - redirecting restored data 1408
 - requirements 1399
 - restoring data 1408
 - setting default options 1412
 - sharing storage on NDMP servers 1402
 - viewing server properties 1413
 - viewing storage device properties 1414
- network
 - configuring options for backup jobs 203
 - overview of backup networks 716
- network and security
 - configuring 718
- network traffic
 - reducing in CAS 1343
- network-attached storage NDMP server
 - backing up 1399
- notification
 - configuring 312
 - configuring group recipient 316
 - editing recipient properties 318
 - enabling 314
 - removing recipient from group 317
 - sending for completed jobs 321
 - setting up 311
 - stopping 322
- notifications
 - overview 303
- NTFS
 - partition 859

O

- off-site storage of backups 132

- offhost backup
 - best practices 1390, 1393
 - configuring 1391
 - host computer
 - defined 1388
 - issues with hardware providers 1396
 - overview 1388
 - setting default options 1385
 - single volume snap 1386, 1392
 - transportable snapshots
 - defined 1388
 - troubleshooting 1394
- one-time conversion to virtual machine 622
- open files
 - unable to back up 858
- OpenStorage devices
 - configure 419
 - data lifecycle management 425
 - disaster recovery of 1002
 - editing properties 422
 - prerequisites 418
 - requirements 972
- Operations Overview report 828
- optimized duplication 997
- Oracle Agent
 - advanced restore options 1253
 - authentication credentials 1240
 - deleting 1242
 - setting 1241
 - authentication credentials options 1242
 - authentication for Oracle operations 1240
 - back up with 1243
 - Backup Exec Agent Utility options 1231
 - backup options 1247
 - configuring 1225
 - DBA-initiated backup 1246
 - DBA-initiated job settings
 - create template for 746
 - DBA-initiated jobs
 - job template name for 1237
 - DBA-initiated restore 1251
 - default options 1227
 - features 1223
 - install 1225
 - Linux servers
 - configuring an Oracle instance 1235
 - deleting an Oracle instance 1238
 - editing an Oracle instance 1238
 - enabling database access 1239

- Oracle Agent *(continued)*
 - Linux servers *(continued)*
 - viewing an Oracle instance 1237
 - multiple data streams
 - specify 1248, 1254
 - Oracle Net Service name 1230
 - port
 - configure for Oracle operations 1240
 - publish Oracle databases on Linux 1236
 - Real Application Cluster (RAC) 1235, 1245
 - recovery catalog 1230, 1237
 - redirected restore 1252
 - restore 1249
 - restore options 1252
 - update credentials for instances 1228, 1236, 1243
 - Windows computers
 - configuring an Oracle instance 1226
 - deleting an Oracle instance 1232
 - editing an Oracle instance 1231
 - enabling database access 1232
 - viewing an Oracle instance 1230
- Overnight Summary report 830
- overwrite protection period
 - changing 503
 - defined 496
 - editing 501

P

- parameter files
 - creating 97
- partition
 - creating for robotic library 486
 - FAT 859
 - NTFS 859
 - removing or reconfiguring 489
- Partition Recovery Utility
 - about 1306
 - finding an archive ID 1307
 - log file location 1294
 - logs
 - about 1293
 - requirements 1306
 - running 1307
 - troubleshooting 1309
- pass phrases 732
- password
 - changing for logon account 765
- patterns in NDMP feature excludes 1406

- performance
 - increase during backups of remote Windows computers 952
- PHYSICAL_ONLY utility 1124
- point in time
 - conversion to virtual machine 620
- point in time log restore option
 - SQL Agent 1146
- pop-up alerts
 - enabling or disabling 324
- Port number
 - changing for Remote Media Agent for Linux 1469
- ports used by Backup Exec
 - default 726
 - listening 728
- post-job command
 - configuring 675
 - for backup jobs 674
- Post-Migration Report 109
- pre-installation checklist 41
- pre-job command
 - configuring 675
 - for backup jobs 674
- pre-upgrade checklist 108
- preferred server configurations for Exchange DAGs
 - about 1171
 - creating 1172
 - deleting 1173
 - designating a default 1173
 - removing as default 1173
- priority
 - about 270
 - changing for scheduled job 270
- Private cloud
 - create cloud instance 379
 - delete cloud instance 382
 - edit cloud instance 382
 - view cloud instance 382
- private cloud-based storage device
 - configure 380
 - requirements 378
- Problem Files report 831
- properties
 - active job 265
- Publish
 - Linux and Unix computers to Backup Exec servers 1429

- publish
 - to Backup Exec servers
 - using Agent for Windows 958

R

- ralus.cfg
 - about, for the Agent for Linux and Unix 1430
 - about, for the Remote Media Agent for Linux computer 1465
 - configuration options 1432
 - editing configuration options in 1432
- Recently Written Media report 832
- recipients
 - configuring groups 316
 - deleting 318
 - editing 318
 - enabling email and text messages 314
 - removing from a group 317
 - stopping notification 322
- reclaiming disk space automatically 352
- Recover This Computer Wizard
 - requirements 916
 - running 916
- recovered jobs
 - setting thresholds 295
- recovered jobs custom error-handling rule 286
- Recovery Ready Validation Summary report 832
- redirected restore
 - Exchange data 1196
- Remote Administrator
 - installing 84
 - installing using the command line 86
 - running 88
- remote computers
 - push-installing 60
- Remote Media Agent for Linux
 - adding to Backup Exec database 1468
 - backing up data 1474
 - changing port number 1469
 - creating a simulated tape library 1476
 - deleting simulated tape library 1479
 - determining server status 1472
 - how it works 1457
 - ICMP ping 1469
 - managing simulated tape libraries from the command line 1480
 - requirements 1458
 - restoring data 1474
 - Simulated Tape Library options 1477

Remote Media Agent for Linux *(continued)*
 simulated tape library properties 1478
 Tape Library Simulator Utility 1475
 Command line switches 1480
 troubleshooting 1481
 uninstalling 1463
 viewing properties 1472
 viewing properties of simulated tape
 libraries 1477

renaming
 tape media labels 512

repair feature 99

replicated catalogs in CAS 1344

reports

 Alert History 808
 Alert History by Backup Exec server 808
 Audit Log 810
 Backup Job Success Rate 811
 Backup Recommendations 811
 Backup Resource Success Rate 812
 Backup Sets by Media Set 812
 Backup Size By Resource 813
 Cloud Storage Summary 813
 copying custom reports 798
 custom
 changing graph options 797
 changing grouping or sorting 796
 previewing 798
 custom, changing filters 794
 custom, creating 789
 Daily Device Utilization 814
 Deduplication Disk and Cloud Deduplication
 Device Summary 815
 Deduplication Summary 816
 deleting 801
 Device Summary 817
 Disk Storage Summary report 817
 editing 800
 Error-Handling Rules 818
 Event Recipients 819
 Failed Backup Jobs 820
 Jobs Summary 821
 list of standard reports 803
 Malware Scan Summary 821
 Managed Backup Exec Servers 822
 Media Audit 824
 Media Errors 825
 Media Required for Recovery 825
 Media Summary 826

reports *(continued)*

 Media Vault Contents 827
 Move Media to Vault 827
 Operations Overview 828
 Overnight Summary 830
 overview 785
 printing 799
 Problem Files 831
 re-running completed report 800
 Recently Written Media 832
 Recovery Ready Validation Summary 832
 Resource Protected Recently 833
 Resource Risk Assessment 834
 Restore Set Details by Resource 835
 Retrieve Media from Vault 836
 Robotic Library Inventory 836
 running 787
 saving 798
 Scheduled Server Workload 837
 scheduling 787
 Scratch Media Availability 838
 setting defaults 801
 Test Run Results 839
 User Role Details 839
 viewing completed 800
 viewing properties 802

requirements

 Agent for Hyper-V 1066
 Backup Exec 45
 Central Admin Server feature 1323
 conversion to virtual machines 597
 Exchange Agent 1161
 NDMP feature 1399
 off-host backup 1389
 Remote Media Agent for Linux 1458
 synthetic backup 1384
 user rights for backup jobs 150

Resource Protected Recently report 833

Resource Risk Assessment report 834

Restore

 creating restore jobs for the Remote Media
 Agent for Linux 1474

restore job

 pre/post commands 674

Restore Set Details by Resource report 835

Restore Wizard 235

restoring

 about restoring data 235
 canceling a restore job 253

- restoring *(continued)*
 - domain controller into an existing domain 249
 - encrypted data 253
 - Exchange data 1196
 - file system data 244
 - from a backup set 242
 - from a completed backup job 242
 - from a server 241
 - from storage device media 243
 - media created with other backup software 253
 - online restore of a Windows computer 246
 - searching for data to restore 237
 - Shadow Copy Components 251
 - SQL master database 1148
 - starting the Restore Wizard 235
 - System State data 247
 - System State to a domain controller 248
 - UEFI system partitions 252
 - utility partitions 252
- restricted encryption keys
 - defined 732
- restricted logon accounts
 - about 761
- retensioning a tape 565
- retired media
 - defined 493
- Retired Servers server group
 - about 153
 - adding servers 156
- retired tape media
 - moving damaged media 517
- Retrieve Media from Vault report 836
- robotic library
 - cleaning slot 485
 - creating partitions 486
 - example configuration 480
 - initializing when Backup Exec services start 484
 - inventory when Backup Exec services start 480
 - reconfiguring partitions 489
 - removing partitions 489
 - setting up hardware 479
 - using with Backup Exec 478
- Robotic Library Inventory report 836
- Role-based access settings
 - add a new user 635
 - configure 634
 - edit role 636
 - role details 637

- RSS Reader
 - customizing 130
 - options 130
 - overview 130
 - removing default RSS feed 130
 - viewing articles 130
- runtime scripts, for Agent for Linux and Unix 1447

S

- SAN
 - hardware errors 862
 - resetting the SAN 863
 - troubleshooting 860
 - troubleshooting offline storage devices 860
- schedule
 - about 216
 - availability window 216
 - configuring for backup jobs 644
 - conflicts 217
 - deleting exclude dates 697
 - exclude dates for all backups 695
 - exclude dates for individual backup jobs 218
 - exporting exclude dates 698
 - importing a list of dates to exclude 695
 - including and excluding dates 217
 - including dates 217
- scheduled jobs
 - about managing and monitoring 261
 - changing priority 270
 - deleting 272
 - holding 266
 - removing a hold 267
 - running immediately 270
 - statuses 289
- Scheduled Server Workload report 837
- scratch media
 - creating 506
 - defined 493
- Scratch Media Availability report 838
- SCSI
 - pass-through mode for tape drives 477
 - setting address for robotic library drives 479
- Search 237
- Search Wizard 237
- seeding a deduplication disk storage device 997
- server groups
 - about 153
 - adding servers to 154
 - backing up 158

- server groups *(continued)*
 - creating 153
 - editing 155
 - hiding 154
 - removing 157
 - removing servers from 155
 - retiring servers 156
 - viewing 154
- Server properties
 - Remote Media Agent for Linux 1472
- server properties
 - about viewing 777
 - viewing 778
- service account
 - changing credentials 771
- services
 - changing service account credentials 771
 - changing startup options 773
 - starting and stopping 771
- SGMon 875
- Shadow Copy Components
 - file system 188
 - restoring 251
- SharePoint Agent
 - about 1204
 - adding a farm 1207
 - backing up SharePoint data 1207
 - disabling or enabling communication between
 - web servers and Backup Exec 1216
 - disaster recovery for SharePoint 1218
 - installing 1205
 - overview 1204
 - requirements 1205
 - restoring SharePoint data 1214
 - system requirements 1205
 - using with SharePoint Server and SharePoint
 - Foundation 1206
- SharePoint farms
 - adding 1207
 - viewing properties 1217
- silent mode installation 90
- Simplified Disaster Recovery
 - Advanced Disk Configuration
 - about 926
 - boot managers 916
 - Contents of the Simplified Disaster Recovery
 - disk image 909
 - editing the default path 890
 - enabling backups for 885
- Simplified Disaster Recovery *(continued)*
 - installing 877
 - Microsoft Exchange Server
 - recovering 924
 - Microsoft Hyper-V hosts 924
 - Microsoft SQL Server
 - recovering 924
 - no internet connection
 - ADK 10 906
 - Windows ADK 903
 - OS/2 boot manager
 - restoring 916
 - overview 877
 - Recover This Computer Wizard
 - encrypted backup sets 916
 - restoring from a remote Backup Exec
 - server 916
 - recovering IBM computers 915
 - recovery requirements in SDR 916
 - requirements 878
 - setting an alternate location 889
 - SharePoint Portal Server
 - recovering 924
 - storage pools and storage spaces
 - recovering 922
 - Windows Server 2012 to Windows Server
 - 2016 897
 - Windows Server 2019 893
 - Windows Server 2022 893
- Simplified Disaster Recovery disk image
 - contents 909
- Simulated tape library
 - creating 1476
 - delete 1479
 - viewing properties 1477
- single block mode
 - setting for tape drives 476
- snapshot technology
 - using with Exchange Agent 1180
- SNMP
 - configuring system service for Windows 328
 - installing WMI provider 329
 - object identifier prefix 325
 - traps
 - defined 325
 - traps for alerts 325
- son media rotation strategy 523
- Sort, filter, and copy 117

- splash screen
 - show at startup 699
- SQL
 - restore to named transaction 1146
- SQL Agent
 - backing up
 - backup methods 1138, 1146
 - consistency check after backup 1136, 1144
 - consistency check recommendations 1123
 - databases 1127
 - strategies for 1121
 - Windows registry 1122
 - consistency check 1124
 - recommendations 1123
 - Database Consistency Check (DBCC)
 - recommendations 1123
 - database snapshots
 - overview 1126
 - disaster recovery 1151
 - manual 1153
 - features 1120
 - installation 1121
 - logon account 1121
 - overview 1119
 - requirements 1121
 - restoring
 - master database 1148
 - point in time log restore option 1146
 - redirecting restores 1146
 - TDE-encrypted database backups 1146
 - very large databases 1146
 - snapshot technology
 - using 1125
 - strategy recommendations 1121
- SQL Server Always On availability groups
 - about 1154
 - adding a listener 1157
 - backing up databases 1158
 - recommendations 1156
 - requirements 1154
 - restoring databases 1159
 - terms 1154
- stages
 - about 220
 - adding 220
 - defined 158
 - editing 221
 - types 221
- stalled jobs
 - setting thresholds 295
- standard view 119
- storage
 - about sharing 557
 - about tape drives and robotic libraries 472
 - backup set retention by storage type 361
 - change to online 559
 - configuring for backup jobs 652
 - deleting 558
 - disabling 564
 - editing global settings 548
 - enabling 564
 - Hot-swappable Device Wizard 472
 - pausing 564
 - renaming 559
 - sharing deduplication devices 995
 - unpausing 564
- storage and media data
 - location of in CAS 1324
- storage device pools
 - about 527
 - adding or removing devices 531
 - changing defaults 530
 - creating 527
 - Least Free Space First device 530
 - Most Free Space First device 530
 - system-defined 527
- storage devices
 - installing 41
- storage operations
 - cataloging 560
 - cleaning a drive 567
 - ejecting media 567
 - exporting expired media 576
 - exporting media 576
 - format WORM 565
 - importing media 568
 - initializing a robotic library 564
 - inventorying 562
 - inventorying and cataloging 563
 - locking the front portal 578
 - overview 534
 - retensioning 565
 - scanning 561
 - unlock front portal 578
- storage pools and storage spaces, recovering with
 - SDR 922
- storage trending 332

- synthetic backup
 - baseline 1383
 - encryption
 - requirements for 1384
 - requirements 1384
- system logon account
 - about 760
 - creating 768
- system requirements
 - Backup Exec 45
- System State
 - restoring 247
 - restoring to a domain controller 248

T

- Tabs 115
- tape drives
 - buffer count 476
 - buffer size 475
 - default settings 476
 - high water count 476
 - statistics 478
- Tape Library Simulator Utility
 - creating a simulated tape library 1476
 - deleting library 1479
 - overview 1475
 - running from command line 1480
 - viewing properties 1477
- tape media
 - associating with a media set or vault 520
 - damaged 517
 - deleting 518
 - scanning barcode labels 517
 - with excessive errors 517
- tape media label
 - barcodes 511
 - imported 511
 - overview 509
 - renaming 512
- tape media operations
 - associating media with media sets 520
- tapes
 - DLT tape drive 853
- technical support
 - contacting 869
- test run job
 - about 226
 - configuring 659
 - running 227

- Test Run Results report 839
- text message notification
 - configuring 312
- transform, creating 75
- Transparent Database Encryption
 - SQL Agent 1146
- tree view 119
- Troubleshooting
 - Remote Media Agent for Linux 1481
- troubleshooting
 - Backup Exec performance
 - improving 867
 - backup issues 858
 - error messages 857
 - hardware-related issues 851
 - installation issues 864
- true image restore
 - overview 1386
- trust
 - establishing 953
 - establishing for a remote Linux and Unix
 - computer 1428
 - establishing for a Remote Media Agent for Linux
 - computer 1465
- typical installation 48

U

- uninstallation
 - Backup Exec 110
 - Backup Exec features from local Backup Exec
 - server 111
 - using command line 111
- Unique Message Identifier (UMI) error code
 - viewing 282
- unrecognized media 491
- updates
 - installing to Agent for Windows 74
 - viewing what is installed 103
- upgrades
 - checklist 108
 - overview 106
- USB tape devices
 - reconnecting 472
- User Role Details report 839
- utility partitions
 - restoring 252

V

- vault
 - scan barcode labels to move tape media 517
- vault rules for media sets 515
- verify operation
 - configuring for backup jobs 661
- Veritas QuickAssist Help Tool 870
- version of Backup Exec, displaying 116
- vhd files
 - about managing 170
- vhdx files
 - with GRT 1068
- Viewing information on the administration console 119
- virtual machine
 - backing up 1013
- virtual machine conversion
 - adding a stage for 613
 - after a backup job 605
 - from point in time 620
 - overview 589
 - requirements 597
 - setting default options 628
 - simultaneous with backup job 599
- virtual machines
 - automatic protection for Hyper-V 1086
- virtual tape library
 - DirectCopy to physical devices 229–230
- virtual-based backup 139
- VMware vCenter Server, adding 1008
- VMware virtual machines
 - push-installing Agent for Windows 1011
- volume level backups
 - automatic exclusion of SQL data 1127
- VSS
 - perform consistency check before Exchange backup 1184, 1191
 - using to protect Exchange data 1180

W

- Windows change journal
 - defaults set by Backup Exec 201
 - resolving errors 201
 - using to determine backed up status 200
- Windows deduplication volume
 - backing up 170
- Windows Management Instrumentation (WMI)
 - adding WMI capability 328

- Windows registry
 - backing up with SQL Agent 1122
- Windows Server 2012
 - Read Only Domain Controller 46
 - Server Core 46
- Windows Server 2012 and later
 - with the Agent for Hyper-V 1068
- Windows user rights 150
- WMI
 - installing performance counter provider 328
 - installing SNMP provider 329
 - uninstalling performance counter provider 329
 - uninstalling SNMP provider 329
- WORM media, about 512