

Veritas NetBackup™ Appliance Capacity Planning and Performance Tuning Guide

Release 3.1 and 3.1.1



Veritas NetBackup™ Appliance Capacity Planning and Performance Tuning Guide

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies 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 Veritas 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 Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies 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. VERITAS TECHNOLOGIES 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 Veritas 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

APPL.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	About this Guide	8
	About this guide	8
	About the intended audience	9
	About the NetBackup appliance documentation	9
Section 1	Capacity planning	13
Chapter 2	Analyzing your backup requirements	14
	Analyzing your backup requirements	14
	What do you want to backup?	14
	How much data do you want to back up?	17
	When should the backup take place?	18
	What is the retention period?	20
	Record your backup requirements	21
Chapter 3	Designing your backup system	31
	Addressing use cases of backup systems for enterprises	31
	Addressing use cases of backup systems for remote or branch offices	34
	About NetBackup appliances	35
	About the Master Server role	39
	About the Media Server role	40
	About NetBackup 53xx High Availability solution	41
	Selecting new appliances	41
	Selecting a media server	41
	Selecting a master server	42
Section 2	Best Practices	44
Chapter 4	About the best practices section	45
	About best practices	45
	References to maintenance sites	45

Chapter 5	About implementing deduplication solutions	46
	About implementing deduplication solutions	46
	General recommendations	46
	Oracle	47
	Microsoft SQL	48
	DB2	49
	Sybase	50
	Lotus Notes	51
Chapter 6	Network consideration	52
	About network considerations	52
	About Fibre Channel connectivity	53
	About SAN zone configurations	53
	About zoning the SAN for NetBackup appliances	54
	Validating network bandwidth	57
Chapter 7	Storage configuration	61
	About storage configuration	61
	About configuring a shared storage pool	66
	Calculating the basic stream count for backups	68
	About moving a storage partition for better performance	68
	Moving a partition	69
	Moving a partition using the NetBackup Appliance Shell Menu	71
	Moving the MSDP partition from a base disk to an expansion disk for optimum performance	72
Chapter 8	Generic best practices	82
	Generic best practices	82
	About Notification settings	83
	About IPMI configuration	84
	Disaster recovery best practices	85
	Job performance	86
	Architecture	87
	NetBackup Catalog Backup	87
	Patching with the SLP (storage lifecycle policies)	88
	VMware backups using appliances	89
	Improving NetBackup Appliance restore performance	90

Section 3	Performance tuning	91
Chapter 9	Role-based Performance Monitoring	92
	Role-based configuration as a factor affecting performance	92
	What affects the performance of a Master Server?	92
	What affects the performance of a media server (MSDP)?	94
Chapter 10	Optimize network settings and improve performance	99
	Optimize network settings and improve performance	99
	SAN Fiber Channel setup	100
	Network bonding	101
	VMware VADP	102
	Jumbo frame implementation for increased MTU	103
Chapter 11	Storage configurations	104
	Storage configurations	104
	Deduplication disk I/O and RAID level settings	105
	RAID Controller operations	106
	RAID Controller commands	108
	Deduplication load balancing	109
	Storage Lifecycle Policies	109
	Auto Image Replication (AIR)	109
	AdvancedDisk settings	110
	Tape Out operations	110
Chapter 12	NetBackup Appliance tuning procedures and performance monitoring	112
	About diagnosing performance issues	112
	About CPU monitoring and tuning	114
	About memory monitoring and tuning	115
	About network monitoring and tuning	116
	About I/O monitoring and tuning	117
	General guidelines to spot a resource bottleneck	121
	About performance tuning practices	121
	About I/O performance tuning	122
	About Oracle backup and restore performance tuning	122
	Setting NET_BUFFER_SZ to 0 on the NetBackup client for better WAN Optimization performance	123
	About tuning procedures and performance monitoring	123

	NetBackup Client performance	124
	Other Performance Monitoring commands	127
Section 4	Quick reference to Capacity planning and Performance tuning	128
Chapter 13	Capacity Planning checklist	129
	Checklist for Capacity planning	129
	Capacity measuring worksheet	137
Chapter 14	Best practices checklist	141
	Best practices checklist	141
Chapter 15	How to monitor the performance of your appliance	147
	About monitoring the performance of your appliance	147
	Performance Monitoring matrix	149
Index		152

About this Guide

This chapter includes the following topics:

- [About this guide](#)
- [About the intended audience](#)
- [About the NetBackup appliance documentation](#)

About this guide

NetBackup appliance is a high-performance data protection application that caters to, large sized enterprises as well as mid sized organizations. It can be scaled to manage large-scale data centers with multiple clients or smaller data centers at remote offices or branch offices. This manual is for administrators who want to analyze, evaluate, and tune the performance of their NetBackup Appliances.

[Table 1-1](#) provides a detailed description of all the sections included in this guide:

Table 1-1 Parts and sections in this guide

Capacity planning	Best practices	Performance tuning
These sections help you to analyze your backup requirements, identify your pain points, and plan on putting together a backup solution that is perfect for your environment	These sections provide a reference for implementing the planned backup solution that you have identified in the Capacity Planning chapter. It lists the ideal settings and scenarios to configure the backup solution.	This section helps to monitor and improve your backup solution to obtain the optimum performance.

Note: The *Capacity Planning and Performance Tuning guide* is an improvement-in-progress guide, that means it will be updated improved each time we learn of new tools, and processes for capacity planning or performance improvements. You can refer to the [NetBackup Appliance Documentation page](#) for the most updated versions of the documentation.

About the intended audience

This guide is intended for the end users that include consultants, partners, system administrators, and IT technicians who are tasked to analyze the backup requirements of their organization to identify a complete backup solution. This guide can also enable the end users to monitor their appliances and inform them about the affect each type of process has on the backup activity.

About the NetBackup appliance documentation

The following documents help to ensure that you can successfully install, configure, and use your appliance. In addition, you can find information about the appliance hardware documents from the following table.

All these documents are posted on the [NetBackup Appliance Documentation page](#).

Table 1-2 NetBackup Appliance Software documentation

Guide	Description
<i>NetBackup™ 52xx Initial Appliance Configuration Guide</i>	This document guides you through the 52xx configuration process from the NetBackup Appliance Web Console or from the NetBackup Appliance Shell Menu.
<i>NetBackup™ 53xx Initial Appliance Configuration Guide</i>	This document guides you through the 53xx configuration process from the NetBackup Appliance Web Console or from the NetBackup Appliance Shell Menu.
<i>NetBackup Appliance Upgrade Guide</i>	This document guides you through the required steps to upgrade a NetBackup appliance.
<i>NetBackup™ Appliance Administrator's Guide</i>	The <i>NetBackup™ Appliance Administrator's Guide</i> contains the following types of information: <ul style="list-style-type: none">■ Deployment information■ Administering your appliance■ Monitoring information

Table 1-2 NetBackup Appliance Software documentation (*continued*)

Guide	Description
<i>NetBackup™ Appliance Command Reference Guide</i>	The <i>NetBackup™ Appliance Command Reference Guide</i> provides a complete list of the commands that are available for you to use through the NetBackup Appliance Shell Menu.
<i>Appliance Management Guide</i>	This document helps you to use the Veritas Appliance Management Console to centrally manage multiple appliances. The Veritas Appliance Management Console provides enterprise-wide monitoring and management of NetBackup appliances. With 3.1 and later, you can manage software upgrades or install EEBs on multiple appliances.
<i>NetBackup Appliance Release Notes</i>	This document contains information about this version of NetBackup Appliance. It contains brief descriptions of new features within the release, operational notes that apply to the release update, and any known issues.
<i>NetBackup Appliance Troubleshooting Guide</i>	<p>This document provides a general overview of how to troubleshoot NetBackup appliance issues and an explanation of the appliance troubleshooting tools and log files.</p> <p>If you need more specific troubleshooting information about a particular issue, go to the NetBackup Appliance page on the Veritas Support website. You can use the search function to look for articles relating to specific issues.</p>
<i>NetBackup Appliance Capacity Planning and Performance Tuning Guide</i>	This document contains information on how to optimize your backup environment and your NetBackup appliance. It helps you to analyze your backup requirements and design a system that best fits your needs.
<i>NetBackup Appliance Fibre Channel Guide</i>	This document describes the supported Fibre Channel (FC) capabilities and configurations for NetBackup appliances.
<i>NetBackup Appliance iSCSI Guide</i>	This document describes how iSCSI works on the NetBackup appliance.
<i>NetBackup Appliance Decommissioning and Reconfiguration Guide</i>	This document describes how to decommission and reconfigure a NetBackup appliance.

Table 1-2 NetBackup Appliance Software documentation (*continued*)

Guide	Description
<i>NetBackup Appliance SNMP Trap Reference Guide</i>	This document provides a complete list of the NetBackup Appliance SNMP traps. It describes what each trap means and the recommended actions for when an error occurs.
<i>NetBackup Copilot for Oracle Configuration Guide</i>	This document outlines how to configure Copilot using NetBackup and the NetBackup Appliance.
<i>NetBackup Appliance Third-party Legal Notices</i>	<p>The <i>NetBackup Appliance Third-party Legal Notices</i> document lists the third-party software that is included in this product, and it contains attributions for the third-party software.</p> <p>This document is available from the following website: https://www.veritas.com/about/legal/license-agreements</p>
<i>NetBackup™ Appliance AutoSupport 2.0 Reference Guide</i>	This document contains the information about the AutoSupport 2.0. It helps you to understand the deployment of the AutoSupport infrastructure, and how does the AutoSupport infrastructure analyzes the Call Home data from each appliances.
<i>NetBackup™ 53xx Appliance High Availability Reference Guide</i>	This document contains the information about the High Availability (HA) solution. It helps you to understand the deployment of the high availability configuration.

Table 1-3 NetBackup Appliance Hardware documentation

Guide	Description
<i>NetBackup™ 5230 Appliance and Storage Shelf Product Description</i>	This guide introduces you to the NetBackup 5230 Appliance and Storage Shelves.
<i>NetBackup™ 5240 Appliance Product Description</i>	This guide introduces you to the NetBackup 5240 Appliance and Storage Shelf.
<i>NetBackup™ 5330 Appliance Product Description</i>	This guide introduces you to the NetBackup 5330 Appliance and Storage Shelves.
<i>NetBackup™ 5340 Appliance Product Description</i>	This guide introduces you to the NetBackup 5340 Appliance and the 5U84 Storage Shelves.

Table 1-3 NetBackup Appliance Hardware documentation (*continued*)

Guide	Description
<i>NetBackup™ Appliance Safety and Maintenance Guide</i>	<p>This document provides safety maintenance information for the following hardware:</p> <ul style="list-style-type: none"> ■ NetBackup 52xx appliances ■ NetBackup 53xx appliances ■ Veritas 3U16 24TB/36TB Storage Shelves ■ Veritas 2U12 49TB Storage Shelf ■ Veritas 5U84 Storage Shelves

See [“About NetBackup appliances”](#) on page 35.

Capacity planning

- [Chapter 2. Analyzing your backup requirements](#)
- [Chapter 3. Designing your backup system](#)

Analyzing your backup requirements

This chapter includes the following topics:

- [Analyzing your backup requirements](#)

Analyzing your backup requirements

Many factors can influence your backup strategy. Analyzing these factors can help you to make better backup decisions based on your site's priorities. This section helps you to analyze the backup requirements of your organization by answering the following core questions:

- [What do you want to backup?](#)

Note: You can use the provided templates in the [Record your backup requirements](#) section to gather backup requirements for your organization.

- [How much data do you want to back up?](#)
- [When should the backup take place?](#)
- [What is the retention period?](#)

What do you want to backup?

The first step in analyzing your backup requirements is to identify the systems that you need to backup. The term **Systems** includes any component that contains vital data which needs to be backed up. There are several parameters and considerations

that you will have to take into account for each type of system. [Table 2-1](#) lists some of these systems and the considerations made to back them up:

Table 2-1 Systems to be backed up

System	Considerations to be made
Computer systems	<div>Identify and record the following information, for each computer system:</div> <div><ul style="list-style-type: none">■ Host name■ Operating system and its version■ Location■ Network technology■ Whether the computer system is directly or remotely connected to the network■ Number of resources to be backed up■ Number of disk drives or tapes, connected to each system■ Model type of each connected drive or library connected■ Applications installed on the computer system that are to be backed up</div> <div>Note: You can use the templates provided in the Template to record computer system information to gather backup requirements for your organization.</div>

Table 2-1 Systems to be backed up (*continued*)

System	Considerations to be made
Databases and Application servers	<p>Identify and record the following information, for each database or application server:</p> <ul style="list-style-type: none">■ Host name■ Location■ Operating system and its version■ Server type and application release numbers■ Network technology - 10 GB Ethernet, 1 GB Ethernet, or SAN Clients■ Database types, database engines, and their version numbers■ Method used to backup these databases■ Whether the databases or application server can be backed up while they are online or offline <p>To backup an online database, you need a NetBackup database agent specific to the database engine. Recording this information can help you to identify the various database agents you may require based on your backup requirements. For the compatibility list of all NetBackup database agents, see http://www.veritas.com/docs/000006088.</p> <ul style="list-style-type: none">■ Whether you want to backup the entire database or the raw partitions. <p>If you decide to take databases backups using raw partitions by Snapshot Client backup, this may result in a backup size equal to the total size of your raw partition. Identifying this information can help you decide how you would like to backup your databases</p> <p>Note: You can use the templates provided in the Template to record database information to gather backup requirements for your organization.</p>
Virtualized datacenters	<p>Identify the following information, about the virtualized infrastructure:</p> <ul style="list-style-type: none">■ Hosts within these virtualized datacenters■ Operating systems and their versions■ Type of data stored■ Type of Hypervisor used■ Network technology used - 10 GB Ethernet, 1 GB Ethernet, or VADP over SAN■ Type of backup performed on these virtualized datacenters- incremental, full

See [“Analyzing your backup requirements”](#) on page 14.

See [“How much data do you want to back up?”](#) on page 17.

See [“When should the backup take place?”](#) on page 18.

See [“What is the retention period?”](#) on page 20.

How much data do you want to back up?

When calculating the total amount of data to be backed up consider the current requirements and the expected growth in the data in the near future. [Table 2-2](#) lists the considerations to be made for calculating the amount of data to be backed up.

Table 2-2 Calculating the amount of data to be backed up

Checks to calculate the amount of data to be backed up	Additional comments
Have you computed the total disk space on each individual system?	<ul style="list-style-type: none">■ Calculate the total disk space for each individual system, along with the space used for databases.■ Identify the total number of files to be backed up on each system.■ Calculate the amount of disk space used by a computer system on their disk drives, shared drives, and so on.■ You can use the calculated total size to analyze the potential growth of your system to ensure that your current backup solution accommodates the growth in backup requirements■ Perform thorough analysis of the data being protected for at-least 30 days to get an understanding of data types, growth trends and change rates using the NetBackup Deployment Utility.
Have you identified and added all the data types included in the backup?	<p>Ensure that you have identified and listed all the data types to be included in the backup. For example, does the data contain text, graphics, databases, and so on. Each data type contributes to a specific amount of data in your total backup requirement.</p> <p>Some of the advantage of identifying the data types to be backed up are listed below:</p> <ul style="list-style-type: none">■ Based on this classification you can identify the amount of data collected under each type, this calculated amount can help you compute the expected growth in data in the near future.■ The other advantage of identifying data types can help you decide upon the relevance and priority of backing up a particular data type.
How compressible is the data?	<p>Identify how much amount of data from your total capacity is compressible or not. Based on this information you can ascertain the amount of space to be allocated to backup data that is compressible and to data that is incompressible.</p>
Will the data be encrypted?	<p>Identify if you want to encrypt the data during backup. Based on the level of security applicable to the data being backed up this decision can be taken.</p> <p>Note: You will also need to consider the additional time encrypted backups may take. Based on which your backup schedules will get affected.</p>

Table 2-2 Calculating the amount of data to be backed up (*continued*)

Checks to calculate the amount of data to be backed up	Additional comments
Have you added the space on mirrored disks only once?	Ensure that you consider the amount of data stored on these repositories only once, if you use storage replication or disk mirroring techniques.
Have you added the size of your database backups to your final calculations?	Ensure that you add the amount of data stored on all the database management systems maintained in your organization.
What are your data protection requirements?	Understand the data protection requirement for the data to be backed up, like the SLA, disaster recovery, and Backup retention and so on.
Have you added the requirements of the test-backup environment?	Create an independent test-backup environment to ensure that your production environment will work with the changed backup solutions. The test-backup environment helps to keep in study the effects of a new backup solution and of the changes made to the hardware and the software infrastructure over a period of time
Have you added any resulting growth factor to your total backup solution?	Identify the annual data growth factor to calculate capacity increase over years.
Have you computed the amount of data you will need to backup in the next six months to a few years from now?	Ensure that you have computed the growth in amount of data in the next six months to a few years. Based on the expected growth you can identify the strategy to expand your existing storage configurations, as required.

See [“When should the backup take place?”](#) on page 18.

See [“Analyzing your backup requirements”](#) on page 14.

See [“When should the backup take place?”](#) on page 18.

See [“What is the retention period?”](#) on page 20.

When should the backup take place?

To correctly size and configure your backup solution, you must decide on the type and frequency of your backups. When considering a backup schedule you need to consider the following checks:

Table 2-3 Schedule a backup

Checks to be performed	Additional comments
Will you perform daily-incremental backups?	The size calculation and deduplication ratio for taking full backups differ from the size calculation and deduplication ratio required for an incremental backup.

Table 2-3 Schedule a backup (*continued*)

Checks to be performed	Additional comments
When will you perform a full backup? <ul style="list-style-type: none">■ Weekly■ Bi-weekly■ Monthly	When you determine this information helps in calculating the storage capacity requirements for over longer period of time. The size calculation and deduplication ratio for taking full backups differ from the size calculation and deduplication ratio required for an incremental backup.
How many jobs will be run per day?	<p>Considering that it takes 2 seconds / 3 seconds per job, it means that approximately ~15000 jobs per backup window (1000 per hour) are run. If the job count exceeds the 2-second or 3-second rule, you may need to consider increasing the backup window or setup a new domain.</p> <p>Also, the number of jobs is limited by the total number of streams that appliance can take and also on the total throughput of the deduplication storage on the appliance.</p>
What is the window of time that is available to complete each backup?	The length of a window dictates several aspects of your backup strategy. For example, you may want a larger window of time to back up multiple, high-capacity servers.
Would you want to configure a deduplication solution to take backups?	<p>In some scenarios you may want to only use AdvancedDisk for backup. An AdvancedDisk is the normal disk backup, performed without any deduplication.</p> <p>Deciding on whether you want to configure a deduplication solution can help you opt for a mix of MSDP and AdvancedDisk.</p>
Would you want to consider using the advanced NetBackup features such as NetBackup accelerator?	NetBackup Accelerator uses a platform and file system independent track log to intelligently detect changed files and send the changed segments from those files to media server. These changed segments are written to a supported storage pool in the NetBackup appliance. Although the primary benefit from NetBackup Accelerator is for performing full backups, it also reduces a subset of costs in running incremental backups.
Would you be adopting NDMP to perform backups?	The NDMP (Network Data Management Protocol) is used along with NetBackup Accelerator to perform backup and recovery. The NetBackup appliance supports NDMP stream handlers for EMC Celerra and NetApp which results in very high deduplication rate. Using Accelerator we can protect shares as mounted file systems. However, this is not a comprehensive backup of filers.
Do you plan to run an Off-site backup or Local backup?	The backup schedule is greatly affected if the backup is run locally or is transferred to an off-site location. The time to transfer a local backup to the off-site location will need to be estimated and considered at the time of scheduling a backup.

Table 2-3 Schedule a backup (*continued*)

Checks to be performed	Additional comments
Is the computer system local or remote?	Determining the total number of systems that are local or remote can help determine the backup cycles.
Do you want to allow the users to run the backups and restores for their computer systems?	This can help to reduce the time it takes to initiate certain operations. However, user-directed operations can also result in higher support costs and the loss of flexibility in storage allocation and usage for backups. User-directed operations can monopolize media and tape drives when you most need them. They can also generate more support calls and training issues while the users become familiar with the new backup system.
Do you plan to back up any system over a network?	Identify the type of networks you plan to use for scheduling backups. Based on the amount of data that you want to back up and the frequency of those backups, you can also consider installing a private network for taking backups. Identifying this information can help you account for the time required to perform the backups over a network and select the best time period when the backup should be taken Note: It is recommended that you use SAN Client and dedicated HBA on the NetBackup client for taking backups.
What is your policy for backup expiration?	Identify the backup expiration time period for each of your policies, thus determining the possible restore options. For example, one of the commonly adopted policies is to expire your incremental backups after 2 weeks or 1 month and your full backups after 6 months. With this policy, you can restore any daily file change from the previous weeks or month and restore data from full backups for the previous 6 months

See [“What is the retention period?”](#) on page 20.

See [“How much data do you want to back up?”](#) on page 17.

See [“What do you want to backup?”](#) on page 14.

See [“Analyzing your backup requirements”](#) on page 14.

What is the retention period?

Now that you have identified what data to backup and how much data is to be backed up, you now need to identify the length of the retention period for your backup. Retention period is the time period during which data is stored and maintained in a backup system. Upon the completion of the retention period the retained data is finally disposed.

Table 2-4 Identifying the length of the retention period

Factors deciding the retention period	Additional comments
What are the standards specific to your organizational for creating a retention policy?	Most organizations lay out specific guidelines on creating retention policies and the time period they would want to hold the backed-up data. These guidelines are considered vital in determining the retention policy for an organization.
What is your Projects requirement to retain the data?	There can be certain critical projects that may require you to hold the data for a longer period of time than your standard retention policy. Identify such projects and the approximate time required to retain the information. This can help you prioritize and allot specific backup storage to projects with such retention period requirements.
What are the user-specific requirements on retaining data? Are their users who want to retain specific data for a longer period?	There can be user requirements that may need you to alter the generic retention policy to accommodate requests to retain data backed up by specific users or user groups.
What are the business requirements to identify the retention time?	There can be business-specific requirements that mandate you to hold specific data for a longer period. For example, market research, findings of a research and development department.
What are the regulatory or legal department laid requirements to retain the data?	The legal departments or the regulatory bodies of a country may have specific guidelines on retaining data. For example, organizations may require to retain the financial data of their organization for over a specific number of years. Such policies affect the length of the retention period.
What is the size of your NetBackup image database?	The retention period is directly proportional to the number of tapes you need and the size of your NetBackup image database. Your NetBackup image database keeps track of all the information on all your disk drives and tapes. The image database size is tightly tied into your retention period and the frequency of your backups.

See [“When should the backup take place?”](#) on page 18.

See [“How much data do you want to back up?”](#) on page 17.

See [“What do you want to backup?”](#) on page 14.

See [“Analyzing your backup requirements”](#) on page 14.

Record your backup requirements

This section provides templates that you can use and customize to record the backup requirements in your organization.

- See [“Template to record computer system information”](#) on page 22.

- See “[Template to record database information](#)” on page 24.
- See “[Template to record application server information](#)” on page 27.

Template to record computer system information

Host name	<host name>
Operating system	<Operating system>
Operating system version	<version #>
Location	<Location>
Installed Applications	<List the applications installed on the computer system that need to be backed up>
Network	<Select from the network types> LAN <input type="checkbox"/> SAN <input type="checkbox"/> WAN <input type="checkbox"/> Others <input type="checkbox"/> <Provide details>

Network connection type

<Select the network connection type and provide its details>

Fibre Channel

☐

<Provide details>

Ethernet

☐

<Provide details>

Others

☐

<Provide details>

Storage

<Which devices are used to store data? >

Disk Drives

☐

Tapes

☐

Others

☐

Storage - Disk Drives

<Provide information about the disk drives used to store data >

Number of Disk Drives connected

Disk Drive Models used

Amount of data stored in each disk drive

Total amount of data stored on all the disk drives

Storage - Tapes *<Provide additional information about the tapes used to store data >*

Number of tapes

Total amount of data stored on the tapes

Template to record database information

Host name *<Host name>*

Operating system *<Operating system>*

Version *<Operating system version>*

Location *< Location >*

Application type *<Name the application type for example, Oracle Database>*

Network *<Select from the network types and provide additional details. >*

LAN

☐

<Provide details.>

SAN

☐

<Provide details.>

WAN

☐

<Provide details.>

Other

☐

<Provide details.>

Network connection type

<Select the network connection type and provide its details. For example, 10GB Ethernet>

Fibre Channel

☐

<Provide details.>

10GB Ethernet

☐

<Provide details.>

1GB Ethernet

☐

<Provide details.>

Storage

<Which devices are used to store data? >

Disk Drives

☐

Tapes

☐

Others

☐

<Provide details.>

Storage - Disk Drives <Provide additional information about the disk drives used to store data? >

Number of Disk Drives connected

Disk Drive Models used

Amount of data stored

Total amount of data stored on the tapes

Storage - Tapes <Provide additional information about the tapes used to store data >

Number of tapes

Total amount of data stored on the tapes

Type of Data <Provide additional information about the type of data to be stored?>

File System

☐

<Provide details.>

SQL

☐

<Provide details.>

Exchange

☐

<Provide details.>

Oracle

☐

<Provide details.>

VMware

☐

<Provide details.>

Template to record application server information

Host name	<host name>
Operating system	<Operating system>
Version	< version>
Location	< Location>
Applications type	<Name the application type for example, Oracle Database>

Network

<Select from the network types and provide additional details. >

LAN

☐

<Provide details.>

SAN

☐

<Provide details.>

WAN

☐

<Provide details.>

Other

☐

<Provide details.>

**Network connection
type**

<Select the network connection type and provide its details. For example, 10GB Ethernet>

Fibre Channel

☐

<Provide details.>

10 GB Ethernet

☐

<Provide details.>

1GB Ethernet

☐

<Provide details.>

Storage

<Which devices are used to store data? >

Disk Drives

☐

Tapes

☐

Other

☐

Storage - Disk Drives

<Provide additional information about the disk drives used to store data? >

Number of Disk Drives connected

Disk Drive Models used

Amount of data stored

Total amount of data stored on the tapes

Storage - Tapes

<Provide additional information about the tapes used to store data? >

Number of tapes

Total amount of data stored on the tapes

Type of Data

<Provide additional information about the type of data to be stored?>

File System

☐

<Provide details.>

SQL

☐

<Provide details.>

Exchange

☐

<Provide details.>

Oracle

☐

<Provide details.>

VMware

☐

<Provide details.>

Designing your backup system

This chapter includes the following topics:

- [Addressing use cases of backup systems for enterprises](#)
- [Addressing use cases of backup systems for remote or branch offices](#)
- [About NetBackup appliances](#)
- [About NetBackup 53xx High Availability solution](#)
- [Selecting new appliances](#)

Addressing use cases of backup systems for enterprises

The next step after you have analyzed your backup requirements is to start identifying the components that together will design your backup system. While designing a complete and efficient backup system you will also need to identify the core pain points that dictate the choices you make in selecting each component. For example, you are required to backup growing number of clients with large data sets, in this case your pain point would be to ensure that the backup solution you consider can account for expandable storage options

Some of the common pain points that you may come across have been listed below:

- Choosing the right hardware, operating system, mandated IT application, and the backup software. Each of these come in multiple variables and increased interoperability challenges like dealing with too many vendors for maintenance.
- Choosing tapes only for long-term retention.

- Choosing the right Disk storage, as each type as its individual pain points:
 - DAS (Direct-attached storage) requires additional hardware and file system maintenance.
 - SAN (Storage area network) brings more complexity and operational costs.
 - NAS (Network-attached storage) taxes the network.
- Choosing the right solutions to backup clients with large data sets. Large data sets means multiple challenges in dealing with:
 - Large databases and application servers
 - Large NAS Filers
 - Slow networks thus requiring to create a dedicated backup network
 - Unsuccessful scheduled full backups due to issues like exceeded backup window.
- Choosing Virtual Machine protection to address concerns over:
 - Resource constrains
 - Dealing the VM Sprawl, when the number of virtual machines (VM) increases over time, simply because of the ease of creating VMs and not because those VMs are absolutely necessary for the business
 - Contrasting deduplication pools

In [Table 3-1](#) we will be evaluating each of the above pain points as a separate use case and identifying the solutions that Veritas NetBackup Appliance offers to address these pain points:

Table 3-1 Use cases with possible solutions

Pain Point	Possible Solutions using Veritas NetBackup Appliance
Dealing with slow Networks	Multiple connectivity options can be considered: <ul style="list-style-type: none"> ■ Client deduplication processing ■ Standard Data Transfer over IP ■ Standard Data Transfer over FC ■ Optimized Transfer over IP
Dealing with clients containing large data sets	The power to scale-out in multiple dimensions using: <ul style="list-style-type: none"> ■ Client Deduplication ■ Media server Deduplication ■ Storage Expansion

Table 3-1 Use cases with possible solutions (*continued*)

Pain Point	Possible Solutions using Veritas NetBackup Appliance
Dealing with issues of backing up clients with large data sets - SAN Client	<p>Protecting clients with large data set - SAN Client</p> <ul style="list-style-type: none"> ■ High speed streaming through multiple Fiber Channel streams ■ No need for off-host systems (low TCO) ■ Use the processing power of Appliances for deduplication ■ Avoid configuring clients as backup servers - conserve resources for critical applications ■ Replace aging physical and virtual tape libraries
Accommodating backups of storage filers like NetApp and EMC	<p>Protecting NetApp and EMC Filers</p> <ul style="list-style-type: none"> ■ Adopting remote NDMP to free up resources for backing up NetApp and EMC filers ■ 10 GbE links for faster streaming, multiple links may be trunked for even larger pipes ■ Flexibility to stream to tape for long-term retention ■ Deduplication and replication offloaded from NetApp and EMC filers
Achieving high deduplication rates	<p>Intelligent deduplication Stream Handlers for high deduplication rates</p> <p>The NetBackupAccelerator is used to perform efficient scan of all files and send only changed segments to the appliance during each full backup. The operations necessary to generate full backup are performed on the appliance and master server in the background. Thus the full backup is deduplicated by using lesser MSDP storage space.</p>
Accommodating traditional file servers	<p>Protecting traditional file servers and NAS using optimized synthetics, accelerator and client-side deduplication</p> <ul style="list-style-type: none"> ■ Highly recommended for any file server ■ Works for any NAS solution through a NAS Client (NFS/CIFS) ■ After the initial full backup, never run full backup again ■ Lowest effect on clients, network, and media server ■ Deduplication may be done at source or NetBackup Media server side

Table 3-1 Use cases with possible solutions (*continued*)

Pain Point	Possible Solutions using Veritas NetBackup Appliance
Adapting to VMware requirements	<p>With NetBackup appliance direct vSphere Backup can be achieved due to the VMware backup host built into NetBackup 52xx Appliances.</p> <ul style="list-style-type: none">■ For all VMs, 100% off-host backups can be taken with Single Appliance■ Appliance software integrates VMware Access Host with NetBackup 52xx Appliances■ All transport methods are supported (SAN, nbd, nbdssl, hotadd)■ On the host that runs the NetBackup Administration Console or the Backup, Archive, and Restore interface, NetBackup must be at 7.5.0.1 or later.■ The appliance must be version 2.5 or later. If the master server is on a separate host (not on the backup host), the master server must use NetBackup 7.5.0.1 or later.■ The VMware policy type must be used. The FlashBackup-Windows policy type is not supported.■ The appliance does not support iSCSI connections. <p>For more information, refer to the 'About the NetBackup appliance as a VMware backup host' section in the <i>NetBackup™ Appliance Administrator's Guide</i>.</p>

Addressing use cases of backup systems for remote or branch offices

The next step after you have analyzed your backup requirements is to start identifying the components that together will design your backup system. While designing a complete and efficient backup system you will also need to identify the core pain points that dictate the choices you make in selecting each component. For example, you are required to backup growing number of clients with large data sets, in this case your pain point would be to ensure that the backup solution you consider can account for expandable storage options

Some of the common pain points that you may come across have been listed below:

- Managing a tape based backup system can pose the following challenges:
 - A backup server and small tape library
 - Shipping tapes off-site for disaster recovery
 - Integration challenges
 - Managed by non-IT staff or less skilled backup staff under guidance from central site
- Managing a disk based backup system can pose the following challenges:

- A backup server and deduplication-disk storage
- Replication to central site
- Integration challenges
- Managed by non-IT staff or less skilled backup staff under guidance from central site
- Managing backups over WAN can pose the following challenges:
 - Source deduplication and backups over WAN can be time consuming.
 - Link unavailability equals to increased risk, thus bringing down the RPO (Recovery Point Objective)
 - Slow rehydrated restores leading to a poor RTO (Recovery Time Objective)

Considering the above pain points of remote or branch office, Veritas NetBackup Appliance provide a solution that includes the following key factors:

- Provides an ideal backup-in-a-box for remote office/branch office
- Includes a built-in deduplication and local restores
- Replicates only unique data to central site via optimized duplication
- Manage as an independent domain (managed by AIR) or as remote media server (managed by SLPs)
- Complete protection (physical and virtual) powered by V-Ray Cloud Gateway
- Fast, policy based Global Disaster Recovery using AIR (Automatic Image Replication)

About NetBackup appliances

NetBackup appliances provide a simplified solution for NetBackup configuration and the daily management of your backup environment. The goal is to provide a solution that eliminates the need to provide dedicated individuals to manage their backup environment.

The appliances are rack-mount servers that run on the Linux operating system. NetBackup Enterprise Server software is already installed and configured to work with the operating system, the disk storage units, and the robotic tape device.

You can determine what role you want to configure the appliance to perform. You can choose to configure a 52xx appliance as follows:

- As a master server appliance
- As a media server for use with an existing master server appliance

- As a media server for use in an existing NetBackup environment

With each of these 52xx configurations, you get the added benefit of internal disk storage.

Note: 52xx refers to 5220 and later models. See the *NetBackup Appliance Release Notes* for details.

A 53xx appliance is configured as a media server by default. You can choose to configure a 53xx appliance as follows:

- As a media server for use with an existing master server appliance
- As a media server for use in an existing NetBackup environment

Note: The 53xx compute node does not have internal disk space available for backups or storage. The space available from the Primary Storage Shelf and up to five Expansion Storage Shelves can be used for backups.

This appliance version allows for easy expansion of existing NetBackup environments that have NetBackup 8.1 and 8.1.1 or greater installed. The appliance also includes its own browser-based interface. This interface is used for local administration of the network, internal disk storage, tape libraries and much more.

NetBackup appliances support the following features:

- Three interfaces for appliance configuration and management:
 - The NetBackup Appliance Web Console is a web-based graphical user interface. It lets you monitor and manage the specific appliance. This interface is compatible with Internet Explorer versions 11.0 and later, and Mozilla Firefox versions 31.0 and later.
 - The Veritas Appliance Management Console is a web-based graphical user-interface that helps you to centrally manage multiple NetBackup appliances from a single interface. The Appliance Management Console lets you upgrade and install EEBs on multiple appliances. You can always log on to a dedicated appliance and use the NetBackup Appliance Web Console for all the other tasks.
- The NetBackup Appliance Shell Menu is a command line driven interface. For a complete description of all appliance commands, refer to the following document:

NetBackup Appliance Command Reference Guide

- Configuration H of the NetBackup 5240 appliance supports iSCSI connections. All configurations of the NetBackup 5340 appliance support iSCSI. See the *NetBackup Appliance iSCSI Guide* for reference.
- Copilot enables Oracle database administrators to work with NetBackup appliance administrators to perform a streamlined backup and restore process of Oracle databases.
- A NetBackup 5340 appliance supports 1 Primary Storage Shelf and up to 3 Expansion Storage Shelves.
For more information on the hardware enhancements, refer to the *NetBackup 5340 Appliance Product Description Guide*.
- The Primary Storage Shelf and the Expansion Storage Shelf in a NetBackup 5330 appliance support 6TB disks, an enhancement over the previous 3TB disks.

Note: Individual storage shelves contain either the 3TB disks or the 6TB disks, but not both.

For more information on the hardware enhancements, refer to the *NetBackup 5330 Appliance Product Description Guide*.

- Starting with NetBackup appliance version 2.7.1, you can use the fully qualified domain name (FQDN) as the appliance host name.
- Backup of VMware virtual machines. NetBackup appliance supports direct backup of VMware virtual machines. The appliance can back up virtual machines without a separate Windows system as backup host.
- Symantec Data Center Security (SDCS) integration. The SDCS agent is installed and configured when you initially configure your appliance. By default, SDCS operates in unmanaged mode and helps secure the appliance using host-based intrusion prevention and detection technology. In managed mode, this agent ensures that the appliance audit logs are sent to an external SDCS server to be validated and verified.
- BMR integration. When the appliance is configured as a master server, you can enable Bare Metal Restore (BMR) from the NetBackup Appliance Web Console.
- IPv4-IPv6 network support. The NetBackup appliances are supported on a dual stack IPv4-IPv6 network. The NetBackup appliance can communicate with, back up, and restore an IPv6 client. You can assign an IPv6 address to an appliance, configure DNS, and routing to include IPv6 based systems. The NetBackup Appliance Web Console can be used to enter information about both IPv4 and IPv6 addresses.

- **ACSL S Support.** This feature facilitates configuration of NetBackup ACS robotics on the NetBackup appliance. The appliance administrator can change the ACSLS entries in the `vm.conf` file on the local appliance.
- **NetBackup SAN Client and Fibre Transport.** SAN Client is a NetBackup optional feature that provides high-speed backups and restores of NetBackup clients. Fibre Transport is the name of the NetBackup high-speed data transport method that is part of the SAN Client feature. The backup and restore traffic occurs over a SAN, and NetBackup server and client administration traffic occurs over the LAN.
- **NetBackup preinstalled.** Simplifies the deployment and integration into an existing NetBackup environment.
- **Tape out option.** The appliance includes a gigabit, dual-port Fibre Channel host bus adapter (HBA).
Multiple FC ports can be used for tape out, as long as they are solely dedicated to the tape out function. For more information, refer to the *Veritas NetBackup Appliance Network Ports Reference Guide*.
- **Hardware component monitoring.** The appliance can monitor key hardware components such as the CPU, disks, memory, power supply modules, and fans. In addition, the appliance provides an optional Call Home feature that allows proactive monitoring and messaging of these NetBackup components.
- **The NetBackup appliances support the core NetBackup software agents.** The NetBackup agents optimize the performance of critical databases and applications.
See the *NetBackup Administrator's Guide Volume I* for more information about the policy types that are supported for each software agent. And for the latest NetBackup appliance compatibility information, refer to the Hardware Compatibility List on the Support website.
www.netbackup.com/compatibility
- **Flexible hardware configuration.** The appliance can be ordered in a variety of configurations to provide the necessary Ethernet ports. Along with the built-in Ethernet ports on the motherboard, expansion cards can be specified to provide additional 1 GB or 10 GB Ethernet ports. Dual-port and quad-port expansion cards are supported.

For more information about hardware configuration, refer to the *NetBackup Hardware Installation Guide* and *NetBackup Appliance and Storage Shelf Product Description* for the appropriate platform.

The following describes how you can incorporate this appliance into your current NetBackup environment:

Replace unsupported media servers	Replace an existing media server that runs on a platform that is not supported in NetBackup 8.1 and 8.1.1.
Add deduplication capability	<ul style="list-style-type: none"> ■ Add the appliance to an existing NetBackup environment or replace an existing media server that does not support deduplication. ■ Configure MSDP partition on the Appliance for deduplication capability.
Use AdvancedDisk for non-deduplicated backups	<ul style="list-style-type: none"> ■ AdvancedDisk can provide faster restore operation but is not space-optimized like MSDP. This is a good solution for backups that include strict tape out schedules. Backups can be expired after duplication to MSDP and space on AdvancedDisk freed up for next day backups.
Add more storage capability	<p>Add storage capability to existing NetBackup 8.1 and 8.1.1 and greater environments.</p> <ul style="list-style-type: none"> ■ Built-in appliance disk storage for 52xx appliances The internal disks can be used for additional backup storage on a 52xx appliance. Note: The 53xx appliance does not have internal disk space available for backups or storage. The space available from the Primary Storage Shelf and Expansion Storage Shelves can be used for storage. ■ Additional external storage The Storage Shelf is an external unit that provides additional disk storage space. You can add up to four of these units to a NetBackup 5220, 5230, or a 5240 appliance. If you need or want to add a Storage Shelf to an existing or an operational NetBackup appliance, your appliance may first require a hardware and/or a memory upgrade. For more information, contact your NetBackup appliance representative about your expansion needs.
Tape backup	The appliance includes a Fibre Channel host bus adapter card for a TLD tape storage device for archive support.

About the Master Server role

A NetBackup 52xx series appliance can be configured as a master server with its own internal disk storage. You configure and use this appliance much like you would use a regular NetBackup master server. You can schedule backups or start a backup manually. Users with the appropriate privileges can perform restores.

Note: The NetBackup 53xx appliance is a media server by default and is not supported for the master server role configuration.

This appliance role provides a simplified administrative interface for the local network, disk, and storage unit management. However, the majority of NetBackup administration such as backup management must be performed through the traditional NetBackup Administration Console.

For complete NetBackup administration information, see the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* and *Volume II*.

See [“About NetBackup appliances”](#) on page 35.

See [“About the Media Server role”](#) on page 40.

About the Media Server role

In this role, a NetBackup 52xx series appliance operates as a media server with its own internal disk storage.

A NetBackup 53xx appliance is a media server by default. The internal storage in a 53xx appliance cannot be used for storing any data or taking any backups. The internal storage is used for storing the operating system, checkpoints, and logs.

If the media server is at version 3.1, the master server must be a NetBackup 52xx with appliance software version 3.1 or later, or a traditional NetBackup master server with version 8.1 or later. If the media server is at 3.1.1, the master server must be a NetBackup 52xx with appliance software version 3.1.1, or a traditional NetBackup master server with version 8.1.1.

Media server appliances use a simplified administrative interface for the local network and for disk storage management. However, the majority of NetBackup administration such as backup management is performed on the master server.

When you perform the initial configuration on the appliance, you specify the associated master server:

- **Specify master server**

If the media server is at version 3.1, the master server must be a NetBackup 52xx with appliance software version 3.1 or later, or a traditional NetBackup master server with version 8.1 or later. If the media server is at 3.1.1, the master server must be a NetBackup 52xx with appliance software version 3.1.1, or a traditional NetBackup master server with version 8.1.1.

See [“About NetBackup appliances”](#) on page 35.

About NetBackup 53xx High Availability solution

Starting with the release version 3.1 of the NetBackup appliance software, Veritas offers a high availability (HA) solution for NetBackup 53xx appliances. The HA solution guarantees the system throughput and operational availability for data protection operations.

The NetBackup 53xx HA solution is a dual-node solution that provides the following benefits:

- Two NetBackup 53xx compute nodes (a node and a partner node) operate in an active/active mode through an HA configuration.
- The two nodes efficiently share the storage workload.
- All stored data is available and accessible through either node.
- All NetBackup jobs can be run on a node while its partner node is serviced or upgraded.
- Other than deploying new systems for the NetBackup 53xx HA solution, the existing systems can also be converted for the HA solution.

For the detailed information about HA, refer to *NetBackup 53xx Appliance High Availability Reference Guide*.

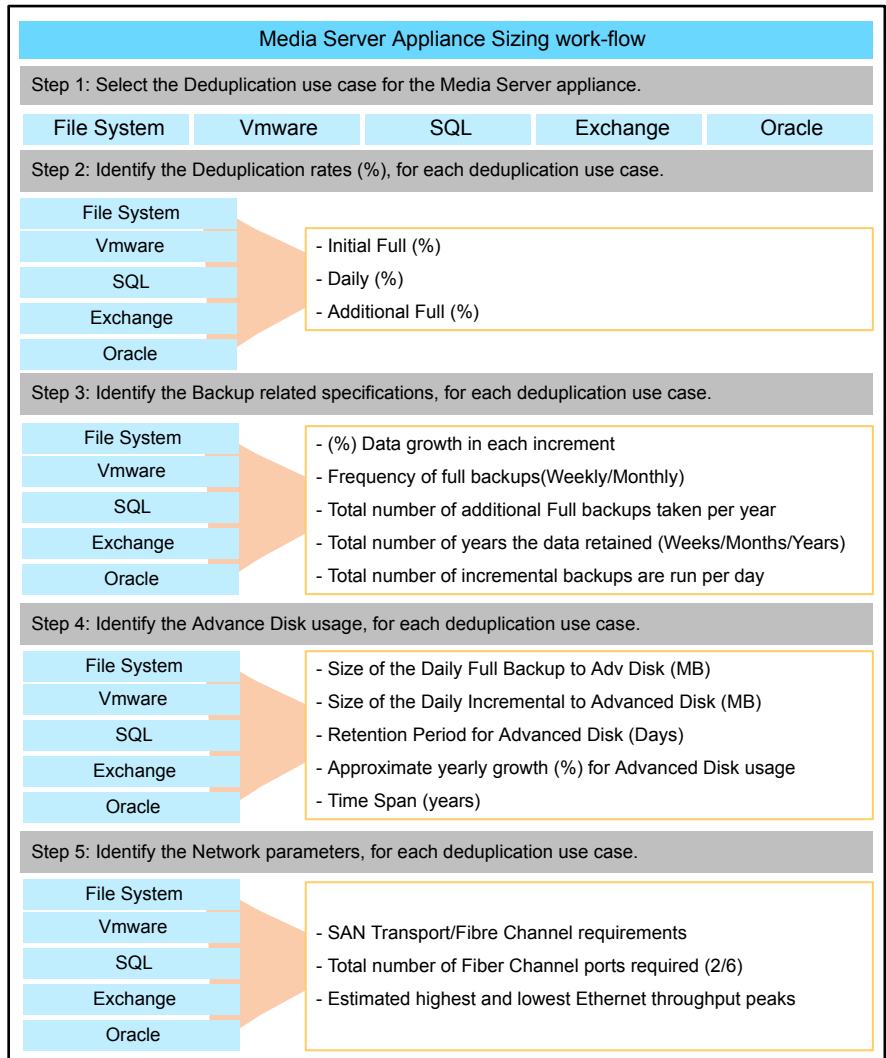
Selecting new appliances

In the previous section we have seen the type of appliances that you can choose from to create your backup solution. This section helps you to select and purchase the components (appliances and storage shelves) based on your requirements.

Selecting a media server

When you plan to buy a media server there are several considerations that you may need to make, so that you can identify the appliance that matches your requirements correctly. [Figure 3-1](#) illustrates the steps you need to follow while selecting a new media server appliance.

Figure 3-1 Selecting a media server

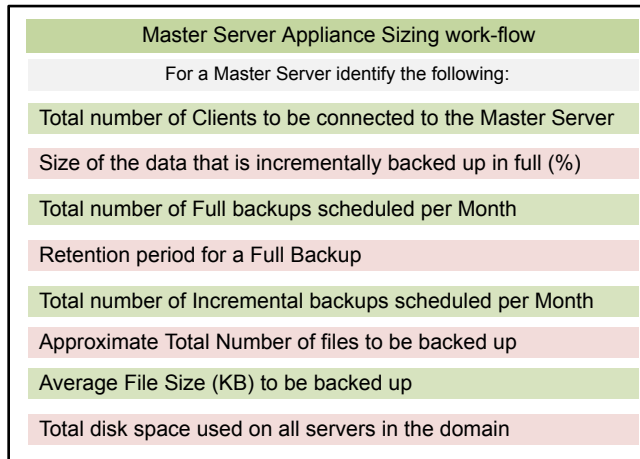


Selecting a master server

When you plan to buy a master server there are several considerations that you may need to make, so that you can identify the appliance that matches your requirements correctly.

Figure 3-2 illustrates the steps you need to follow while selecting a new master server appliance.

Figure 3-2 Selecting a master server



Best Practices

- [Chapter 4. About the best practices section](#)
- [Chapter 5. About implementing deduplication solutions](#)
- [Chapter 6. Network consideration](#)
- [Chapter 7. Storage configuration](#)
- [Chapter 8. Generic best practices](#)

About the best practices section

This chapter includes the following topics:

- [About best practices](#)
- [References to maintenance sites](#)

About best practices

This section lists the best practices for working with the appliance hardware and software. It includes the following sections:

- See “[About implementing deduplication solutions](#)” on page 46.
- See “[About network considerations](#)” on page 52.
- See “[About configuring a shared storage pool](#)” on page 66.

References to maintenance sites

It is always a best practice to ensure that your Veritas NetBackup Appliance is up-to-date with the latest updates. For any recent updates, hot fixes, patches, and other technical advisements, you can use the following links:

- NetBackup Appliance Support website
- <http://www.veritas.com/docs/000010279>
- <http://www.veritas.com/docs/000011514>
- <https://sort.veritas.com/netbackup>

About implementing deduplication solutions

This chapter includes the following topics:

- [About implementing deduplication solutions](#)
- [General recommendations](#)
- [Oracle](#)
- [Microsoft SQL](#)
- [DB2](#)
- [Sybase](#)
- [Lotus Notes](#)

About implementing deduplication solutions

This section lists the recommendations for implementing the Media server-based deduplication (MSDP) and client side deduplication solutions for various databases and applications.

General recommendations

We begin by taking a look at some general best practices that should be considered when you implement a deduplication solution:

- Implement NetBackup database or application agents wherever possible to take advantage of database or application specific stream handlers.

- Database dumps typically result in lower deduplication rates and are not recommended when agents and/or stream handlers are available to perform the task.
- Enable compression at the media server level to achieve higher deduplication rates for databases or applications.
- Leverage client direct backups (client side deduplication) only when there are enough CPU cycles available on the database or application server.
- Use Accelerator and VRay technology wherever possible to further minimize the traffic effect on the network.
- There can be a discrepancy in the MSDP value displayed by the NetBackup appliance software and the actual backup space available or used in the MSDP partition. This discrepancy is caused due to the additional utilization of the MSDP reserve of around 4.4% storage space. The reserved space helps to prevent file system from filling up and affect its performance.

See [“About implementing deduplication solutions”](#) on page 46.

Oracle

To achieve higher deduplication rates while performing Oracle database backups, implement one of the following options:

Table 5-1 Best practices to perform Oracle database backups

Stream based backups (NetBackup for Oracle Agent with RMAN)	<ul style="list-style-type: none"> ■ Implement FILESPERSET = 1 in RMAN backup script for ASM-based backups ■ FILESPERSET can be increased to higher number depending on type of database and change rate. Be cautious and gradually increase to validate the performance improvement and deduplication ratios. Note: FILESPERSET = 4 or FILESPERSET = 5 seems to work well in most scenarios with a little reduction in deduplication ratios. ■ Turn off Oracle compression while performing backups (this is OFF by default). ■ Increase number of channels for faster backups and restores. Note: Setting FILESPERSET > 1 increases performance but reduces dedupe ratios hence Veritas recommends that you begin by experimenting with the different levels of FILESPERSET and identify the most suitable option.
---	---

Table 5-1 Best practices to perform Oracle database backups (*continued*)

Proxy based backups (NetBackup for Oracle Agent without snapshots)	<ul style="list-style-type: none"> ■ Leverage proxy method to enable data movement by NetBackup. ■ Turn off Oracle compression while performing backups (this is OFF by default). ■ Increase number of channels for faster backups and restore.
Proxy based backups with Snapshot Client (NetBackup for Oracle Agent)	<ul style="list-style-type: none"> ■ Leverage proxy method to enable data movement by NetBackup. ■ Enables faster recovery with Instant Recovery using latest snapshots.
Oracle Dump backups (No NetBackup for Oracle Agent)	Leverage IMAGE MODE while dumping to local disk for higher deduplication rates while performing dump backups with NetBackup standard client

See [“About implementing deduplication solutions”](#) on page 46.

Microsoft SQL

To achieve higher deduplication rates while performing SQL server backups, implement one of the following options:

Table 5-2 Best practices to perform SQL server backups

NetBackup for SQL Agent	Use the NetBackup for SQL Agent for better deduplication rates.
<code>pd.conf</code>	Ensure COMPRESSION is set to 1 and MATCH_PDRO = 0 in <code>pd.conf</code> file on the media server for better deduplication rates.
NO_COMPRESSION	Make sure NO_COMPRESSION setting is enabled on SQL Server 2008 and higher.
BLOCKSIZE	Use default BLOCKSIZE for SQL database backups.
BUFFERCOUNT	BUFFERCOUNT defines number of I/O buffers leveraged for a backup process. Use 2 buffers per stripe for good deduplication performance
MAXTRANSFERSIZE	Leverage this parameter for faster backup speeds. Be careful, that you have enough memory while defining both BUFFERCOUNT and MAXTRANSFERSIZE. Otherwise, the backups could result in error 11's and 13's.

Table 5-2 Best practices to perform SQL server backups (*continued*)

Defragmentation	Defragmentation of the SQL databases will also affect the deduplication rates. With SQL databases, more frequently defragmentation is performed, lower the deduplication rates. Instead of performing defragmentation regularly, only perform when certain thresholds are reached.
STRIPES	This parameter determines the number of simultaneous backup streams per single backup operation. Use multiple stripes for faster backups. STRIPES parameter with SQL backups does not seem to have any negative effect on the deduplication rates.

See [“About implementing deduplication solutions”](#) on page 46.

DB2

To achieve higher deduplication rates while performing DB2 data server backups, implement one of the following options:

Table 5-3 Best practices to perform DB2 data server backups

NetBackup for DB2 agent	Use NetBackup for DB2 agent for better deduplication rates
<code>pd.conf</code>	Make sure the <code>pd.conf</code> file has <code>COMPRESSION</code> set to 1 on the media server and <code>SEGKSIZE = 128</code> .
COMPRESS	Do not use <code>COMPRESS</code> parameter with db2 database backup script.
DEDUP_DEVICE	Leverage DB2's new parameter “ <code>DEDUP_DEVICE</code> ” which basically allows for DB2 backup command to identify that the data is being sent to a deduplication device hence resulting in better deduplication rates. This requires DB2 database version to be at least 9.5 Fix Pack 8 and above.
Defragmentation	Database fragmentation can negatively effect the deduplication rates. It is highly recommended to verify that DB2 database is defragmented.

Table 5-3 Best practices to perform DB2 data server backups (*continued*)

Data change rates	<p>Based on the data change rates in the DB2 database, perform the following:</p> <ul style="list-style-type: none">■ Minimal data change rate Limit the number of parallel streams to just 1 for better deduplication rates. Basically, this would be PARALLELISM parameter in DB2 database backup script. This enables the number of tablespaces to be read in parallel by a backup utility to just 1. This will affect the overall backup performance but will achieve better deduplication rates. Increase the number of sessions to generate multiple backup jobs per database instance to achieve higher backup speeds.■ High data change rate requiring high backup speeds: Consider the following for better deduplication rates with high data change rates. Align number of containers in a DB2 database with the number of open sessions in the db2 database backup script to achieve higher deduplication rates. This option needs to be tested on a test/dev environment prior to implementing in production. This would also require additional assistance from the DB2 administrator to understand the number of containers defined in the database.
-------------------	--

See [“About implementing deduplication solutions”](#) on page 46.

Sybase

To achieve higher deduplication rates while performing Sybase database backups, implement one of the following options:

Table 5-4 Best practices to perform Sybase database backups

NetBackup for Sybase agent	Use NetBackup for Sybase agent for better deduplication rates.
pd.conf	Make sure the <code>pd.conf</code> file has COMPRESSION value set to 1 on the media server and SEGKSIZE set to a value of 128.
COMPRESS	Make sure that COMPRESS parameter is not used in the Sybase backup scripts. Otherwise, the deduplication rates are going to be very low.

Table 5-4 Best practices to perform Sybase database backups (*continued*)

STRIPES	Number of stripes defined in a Sybase backup script doesn't seem to impact the overall deduplication performance.
---------	---

See [“About implementing deduplication solutions”](#) on page 46.

Lotus Notes

To achieve higher deduplication rates while performing Lotus notes database backups, implement one of the following options:

Table 5-5 Best practices to perform Lotus notes data server backups

NetBackup for Lotus Notes agent	Use NetBackup for Lotus Notes agent for better deduplication rates
pd.conf	<p>Make sure the <code>pd.conf</code> file has the following settings</p> <ul style="list-style-type: none"> ■ COMPRESSION set to 1 ■ MATCH_PDRO = 0 ■ LOCAL_SETTINGS = 0 on the media server ■ SEGKSIZE = 128
COMPRESSION	<p>Make sure that COMPRESSION is turned off for the following:</p> <ul style="list-style-type: none"> ■ Use LZ1 compression for attachments is disabled on all mailboxes and at the database level ■ Use Domino Attachment and Object Service is disabled on all mailboxes and at the database level ■ Compress database design is disabled on all mailboxes and at the database level ■ Not using a 3rd party compression application for attachments. ■ Compress document data is disabled on all mailboxes and at the database level
Compaction	Make sure that compaction is run at a scheduled interval instead of user initiated compaction (weekly if possible).

See [“About implementing deduplication solutions”](#) on page 46.

Network consideration

This chapter includes the following topics:

- [About network considerations](#)
- [About Fibre Channel connectivity](#)
- [About SAN zone configurations](#)
- [Validating network bandwidth](#)

About network considerations

This section provides generic guidelines and best practices to configure the network settings of your appliance in a given architecture. It provides the hardware consideration and the best possible settings to be considered for configuring your appliance network.

- Enable SAN Client & Fibre Transport (FT) mode data transfer with Network Transport method for improved efficiency of backups.
- Use network bonding for improved throughput and backup performance.
- Use Resilient Network feature of NetBackup to overcome remote office backup bandwidth issues related to latency and to connection intermittence.
- Use 10GbE network infrastructure end-to-end wherever possible for maximum performance.
- Use client-side deduplication to alleviate network congestion, increase concurrency over a given network bandwidth.
- Use WAN optimization settings to address remote office or WAN backup requirements.

Along with above considerations, this chapter includes the following sections:

- See [“About Fibre Channel connectivity”](#) on page 53.
- See [“About SAN zone configurations”](#) on page 53.
- See [“Validating network bandwidth”](#) on page 57.

About Fibre Channel connectivity

The NetBackup appliance supports multiple data transfer features that use Fibre Channel (FC) connectivity to transfer data at high bandwidth:

Table 6-1 Data transfer features that use FC

Data transfer feature	Description
SAN Client	This feature provides high-speed backups and restores of NetBackup clients. The backup and restore traffic occurs over FC, and the NetBackup server and client administration traffic occurs over the LAN.
NetBackup for VMware	This feature provides backup and restore of the VMware virtual machines that run on VMware ESX servers. When using this feature with the SAN Transport backup method, SAN LUNs hosting the VMware datastore should be visible to the media server appliance by the FC ports. Specific ports are not allocated for this function – any port not allocated to another function may be used.
Tape Out	NetBackup 52xx and 53xx appliances support FC connections to TLD tape storage devices allowing for backup data to be written to tape. Appliance software version 2.5 and later includes official support for ACSLS. Connections to tape devices may be direct, or by an FC Switch.

See [“About network considerations”](#) on page 52.

About SAN zone configurations

When configuring SAN for your NetBackup appliance it is important that the network is designed using the SAN zones recommended by Veritas. These SAN zones include **Fibre Transport (FT) Backup Zones** and **Fibre Transport (FT) Storage Zones**. SAN zones are typically configured on the SAN switch within the SAN.

See [“About network considerations”](#) on page 52.

About zoning the SAN for NetBackup appliances

Before you can configure and use the NetBackup Fibre Transport (FT) mechanism, the SAN must be configured and operational.

The NetBackup appliance supports the following SAN configurations:

- Node port (N_Port) switched configuration.
- Fibre Channel arbitrated loop (FC-AL) configuration.
FC-AL hubs are not supported.

For SAN switched configurations, proper zoning prevents Fibre Transport traffic from using the bandwidth that may be required for other SAN activity. Proper zoning also limits the devices that the host bus adapter (HBA) ports discover; the ports should detect the other ports in their zone only. Without zoning, each HBA port detects all HBA ports from all hosts on the SAN. The potentially large number of devices may exceed the number that the operating system supports.

Instructions for how to configure and manage a SAN are beyond the scope of the NetBackup documentation. However, the following recommendations may help you optimize your SAN traffic.

[Table 6-2](#) describes the best practices for zoning the SAN on NetBackup appliances.

Table 6-2 Best practices for zoning the SAN on NetBackup appliances

Guideline	Description
One initiator per zone, multiple targets acceptable.	<p>Veritas recommends that you create zones with only a single initiator per zone. Multiple targets in a single zone are acceptable, only if all of the targets are similar.</p> <p>Note: For data duplication between two NetBackup 52xx or 53xx appliances, you must create zones with only one single initiator and one single target per zone.</p> <p>Tape target resources should be in separate zones from disk target resources, regardless of initiator. However, both sets of resources may share the same initiator.</p>
Be aware of performance degradation when a port is configured for multiple zones.	<p>If you use a single port as an initiator or a target for multiple zones, this port can become a bottleneck for the overall performance of the system. You must analyze the aggregate required throughput of any part of the system and optimize the traffic flow as necessary.</p>

Table 6-2

Best practices for zoning the SAN on NetBackup appliances

(continued)

Guideline	Description
For fault tolerance, spread connectivity across HBA cards and not ports.	To ensure the availability of system connections, if you incorporate a multi-path approach to common resources, pair ports on separate cards for like zoning. This configuration helps you avoid the loss of all paths to a resource in the event of a card failure.
Zone the SAN based on WWN to facilitate zone migrations, if devices change ports.	It is recommended that you perform SAN zoning based on WWN. If switch port configurations or cabling architectures need to change, the zoning does not have to be recreated.

Table 6-3 describes the zones you should use for your SAN traffic.

Diagrams that show the zones are available at the following link.

Note: You must use physical port ID or World Wide Port Name (WWPN) when you specify the HBA ports on NetBackup appliances.

Table 6-3 Appliance zones

Zone	Description
Fibre Transport backup zone	<p>A Fibre Transport backup zone should include only the Fibre Transport traffic between the SAN clients and the appliance.</p> <p>The backup zone should include the following HBA ports:</p> <ul style="list-style-type: none"> ■ The target port of the HBA—connect this port to a Fibre Channel switch port. If you have two HBAs, you can use both of them. The use of two ports provides redundancy. <p>Note: The supplied QLogic FC HBA card in a NetBackup appliance uses a special NetBackup target mode driver for the target port. The target mode driver replaces the default, initiator mode Fibre Channel driver. The target mode driver applies only to the supplied QLogic HBA card.</p> <p>You must define the appliance target port by physical port ID or World Wide Port Name (WWPN). The target mode driver WWPNs are not unique because they are derived from the Fibre Channel HBA WWPN.</p> <ul style="list-style-type: none"> ■ Ports on the SAN client HBAs that connect to the appliance—connect each SAN client HBA port to ports on the same Fibre Channel switch. <p>You can define SAN client ports by either port ID or WWPN. However, if you use one method for all devices, zone definition and management is easier.</p> <p>The ports on the SAN clients use the standard initiator mode driver.</p> <p>To promote multistream throughput, each SAN client should detect all target mode devices of the appliance HBA port or ports in the zone. Each appliance HBA target port exposes two target mode devices.</p> <ul style="list-style-type: none"> ■ Define the zones on the switch so that the client ports and the HBA target ports are in the same zone. <p>Some Veritas appliance models include one or more Fibre Channel HBA cards that can be used for Fibre Transport. If your appliance does not include any of these cards, an authorized Veritas representative must install and configure an approved FC HBA card.</p>

Table 6-3 Appliance zones (*continued*)

Zone	Description
Fibre Transport storage zone	<p>A Fibre Transport storage zone carries the Fibre Transport traffic from a storage source to a storage destination.</p> <p>The traffic can be either for duplication or for backups. For duplication, the deduplicated data is sent to the destination for storage. For backups, the data first travels to the NetBackup 52xx or 53xx appliance and is then sent to the NetBackup deduplication appliance for storage.</p> <p>The storage zone should include the following HBA ports:</p> <ul style="list-style-type: none"> ■ The initiator port of the HBA in the NetBackup appliance—connect this port to a Fibre Channel switch port. It does not have to be the same switch as the backup zone. The NetBackup 52xx or 53xx appliance is the source for the duplication. The initiator ports use the standard initiator mode driver. ■ The NetBackup deduplication appliance ports. For the 5020, connect the target ports (Port 1) of the HBAs in slots 1, 2, and / or 3 to the same Fibre Channel switch. For the 5030, connect the target ports (Port 1) of the HBAs in slots 2 or 5 to the same Fibre Channel switch. The 5020 or 5030 deduplication appliance is the target of the duplication jobs. <p>Note: To use Fibre Channel on a NetBackup 5020 or 5030, you must enable the Fibre Channel communication feature. For details, see the NetBackup Deduplication Appliance Software Administrator's Guide.</p> <ul style="list-style-type: none"> ■ Define the zones on the switch so that the NetBackup appliance initiator port and the NetBackup deduplication appliance target port are in the same zone.
External tape storage zone	<p>If you use a tape library as storage, create a separate zone for that traffic. The tape storage zone does not use NetBackup Fibre Transport; it uses the standard initiator mode driver.</p> <p>The tape storage zone should include a port or ports on the FC HBA in slot 3 of a 52xx appliance, or on the FC HBA in slot 2 or 3 of a 53xx appliance.</p> <p>If slots 5 and 6 do not have FTMS enabled, they can be used as FC ports for tape devices.</p>

Validating network bandwidth

It is recommended to test your network bandwidth before deployment just to ensure that it does not contain any bandwidth bottlenecks. Making sure that the network is performing well prevents any unexpected surprises when performance testing replications, backups and restores with appliances. A variety of tools can be used to test network bandwidth.

You can use the `Nbperfchk` command in all appliances that can measure network and disk read/write speeds as the I/O passes through NetBackup and NetBackup

Appliances. It can be used at pre-deployment stages to measure network speeds before initiating AIR replication between master server domains, or measuring disk write speeds on Appliance storage to check for any storage I/O performance problems before performing a backup and restore demonstration to the customer. Be sure to configure the appliance's network configuration and storage configuration before running any network or storage performance tests. Nbperfchk can be run from the Appliance's Shell Menu: Support->Nbperfchk

To use nbperfchk for network bandwidth tests between two appliances, run nbperfchk as a reader on one appliance and a writer on another, example:

Reader

```
symmaster-a.Support> Nbperfchk run  
Please enter options:  
nbperfchk -i tcp::5000 -o null
```

Writer

```
Symmaster-b.Support> Nbperfchk run  
Please enter options:  
nbperfchk -i zero: -o tcp:symmaster-a:5000
```

In the above example, the symmaster-b appliance is sending data to symmaster-a's 5000 TCP/IP port. The commands provide the following outputs:

Output for the reader

```

symmaster-a.Support > Nbperfchk run
Please enter options:
nbperfchk -i tcp::5000 -o null:
Statistics log are recorded in nbperfchk_results.log
current rcv buff: 262144, set to 524288
current snd buff: 262144, set to 524288
final receive size 262144, send size 262144
    226 MB @ 113.1 MB/sec, 226 MB @ 113.1 MB/sec, at 1369103091
    566 MB @ 113.1 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103094
    905 MB @ 113.1 MB/sec, 339 MB @ 113.1 MB/sec, at 1369103097
   1245 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103100
   1585 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103103
   1925 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103106
   2265 MB @ 113.1 MB/sec, 339 MB @ 113.1 MB/sec, at 1369103109
   2604 MB @ 113.1 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103112
   2944 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103115
   3284 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103118
   3624 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103121
   3964 MB @ 113.2 MB/sec, 339 MB @ 113.1 MB/sec, at 1369103124
   4303 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103127
   4643 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103130
   4983 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103133
   5322 MB @ 113.2 MB/sec, 339 MB @ 113.1 MB/sec, at 1369103136

```

Output for the writer

```

symmaster-b.Support > Nbperfchk run
Please enter options:
nbperfchk -i zero: -o tcp:symmaster-a:5000
Statistics log are recorded in nbperfchk_results.log
current rcv buff: 262144, set to 524288
current snd buff: 262144, set to 524288
final receive size 1048576, send size 262144
    340 MB @ 113.3 MB/sec, 340 MB @ 113.3 MB/sec, at 1369103067
    680 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103070
    1020 MB @ 113.2 MB/sec, 340 MB @ 113.2 MB/sec, at 1369103073
    1360 MB @ 113.3 MB/sec, 340 MB @ 113.4 MB/sec, at 1369103076
    1701 MB @ 113.4 MB/sec, 341 MB @ 113.7 MB/sec, at 1369103079
    2040 MB @ 113.3 MB/sec, 339 MB @ 112.9 MB/sec, at 1369103082
    2381 MB @ 113.3 MB/sec, 340 MB @ 113.3 MB/sec, at 1369103085
    2721 MB @ 113.3 MB/sec, 340 MB @ 113.3 MB/sec, at 1369103088
    3060 MB @ 113.2 MB/sec, 339 MB @ 112.9 MB/sec, at 1369103091
    3400 MB @ 113.2 MB/sec, 340 MB @ 113.3 MB/sec, at 1369103094
    3740 MB @ 113.2 MB/sec, 340 MB @ 113.3 MB/sec, at 1369103097
    4080 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103100
    4420 MB @ 113.2 MB/sec, 339 MB @ 113.2 MB/sec, at 1369103103
    4759 MB @ 113.2 MB/sec, 339 MB @ 113.1 MB/sec, at 1369103106
    5099 MB @ 113.2 MB/sec, 340 MB @ 113.2 MB/sec, at 1369103109
output: Connection reset by peer

```

In the above example, the average network throughput is 113.2 MB/sec. If this network throughput is insufficient for the amount of data being protected, the network infrastructure needs to be examined and additional bandwidth needs to be added.

See [“About network considerations”](#) on page 52.

Storage configuration

This chapter includes the following topics:

- [About storage configuration](#)
- [About configuring a shared storage pool](#)
- [About moving a storage partition for better performance](#)

About storage configuration

The NetBackup Appliance Web Console enables you to manage the storage configuration. You can use the **Manage > Storage** pane to manage the storage space.

The NetBackup 52xx appliances are available for use with up to four storage shelves. The storage shelves provide you with additional disk storage space. After you have physically connected the storage shelves, use the NetBackup Appliance Web Console to manage the storage space.

The NetBackup 53xx appliance must be connected to one Primary Storage Shelf. The storage space can be expanded by using up to five Expansion Storage Shelves. After you have physically connected the Expansion Storage Shelf, use the NetBackup Appliance Web Console to manage the storage space.

Note: The 53xx appliance (base unit) does not have internal disk space available for backups or storage. It only stores the OS, logs, checkpoints etc. The space available from the Primary Storage Shelf and the Expansion Storage Shelf can be used for backups.

If you have a NetBackup 5330 appliance or a NetBackup 5340 appliance with an Expansion Storage Shelf, the following restrictions apply:

- Moving an Expansion Storage Shelf or disks from one 5330 appliance to another 5330 appliance is not supported.
- Moving an Expansion Storage Shelf or disks from one 5340 appliance to another 5340 appliance is not supported.
- Moving an Expansion Storage Shelf or disks from one 5330 appliance to a 5340 appliance is not supported.
- Moving an Expansion Storage Shelf or disks from one 5340 appliance to a 5330 appliance is not supported.
- Moving disk drives within an Expansion Storage Shelf is not supported.

Figure 7-1 provides a bird's-eye view of how storage space is configured within your 52xx appliance.

Figure 7-2 provides a bird's-eye view of how storage space is configured within your 5330 appliance.

Figure 7-3 provides a bird's-eye view of how storage space is configured within your 5340 appliance.

Figure 7-1 NetBackup 52xx Appliance storage space

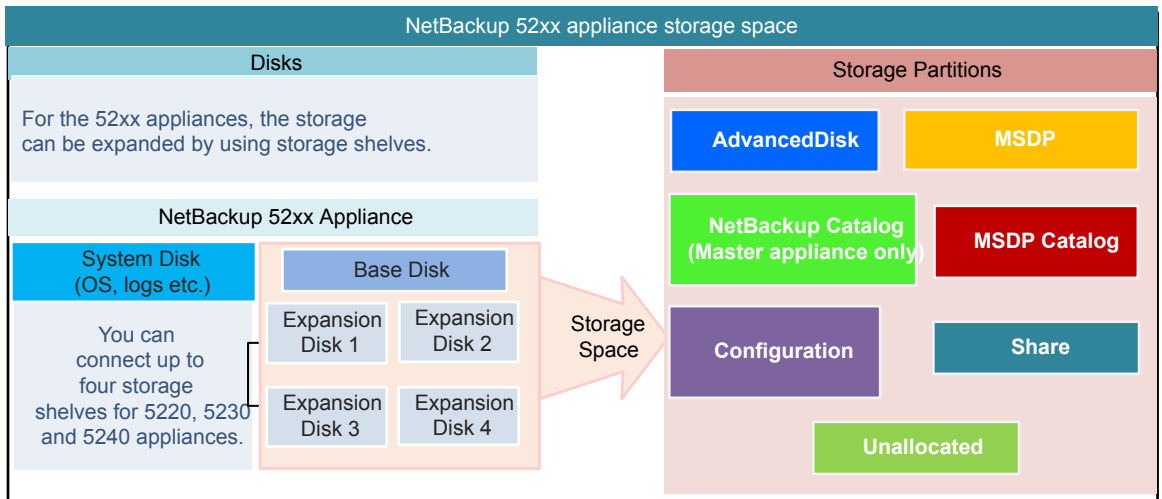


Figure 7-2 NetBackup 5330 Appliance storage space

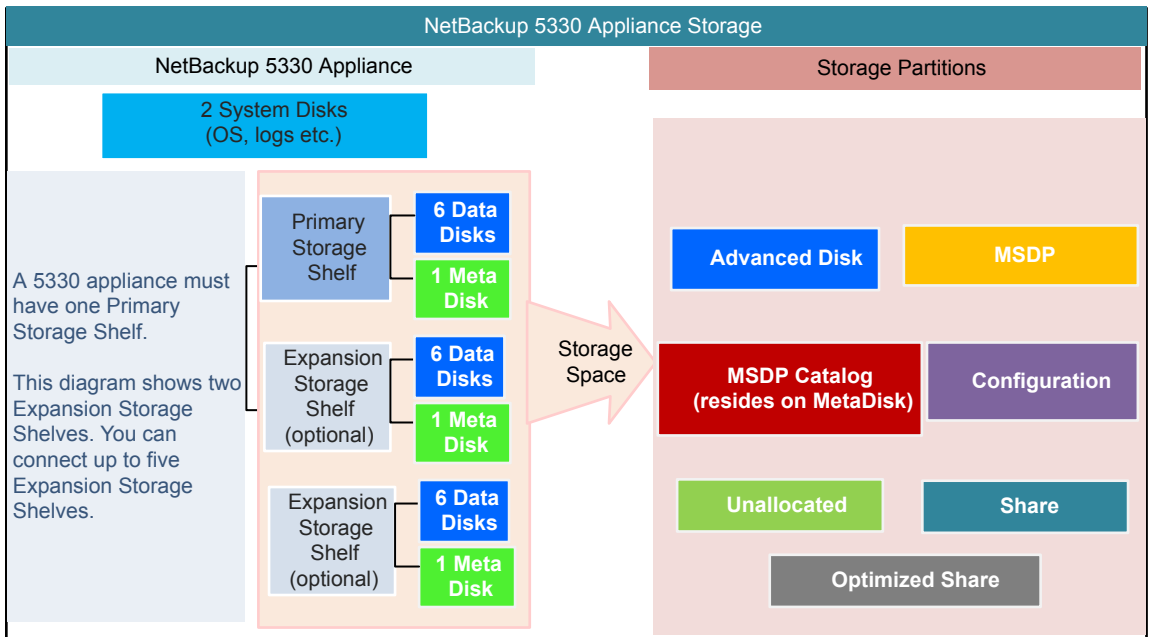


Figure 7-3 NetBackup 5340 Appliance storage space

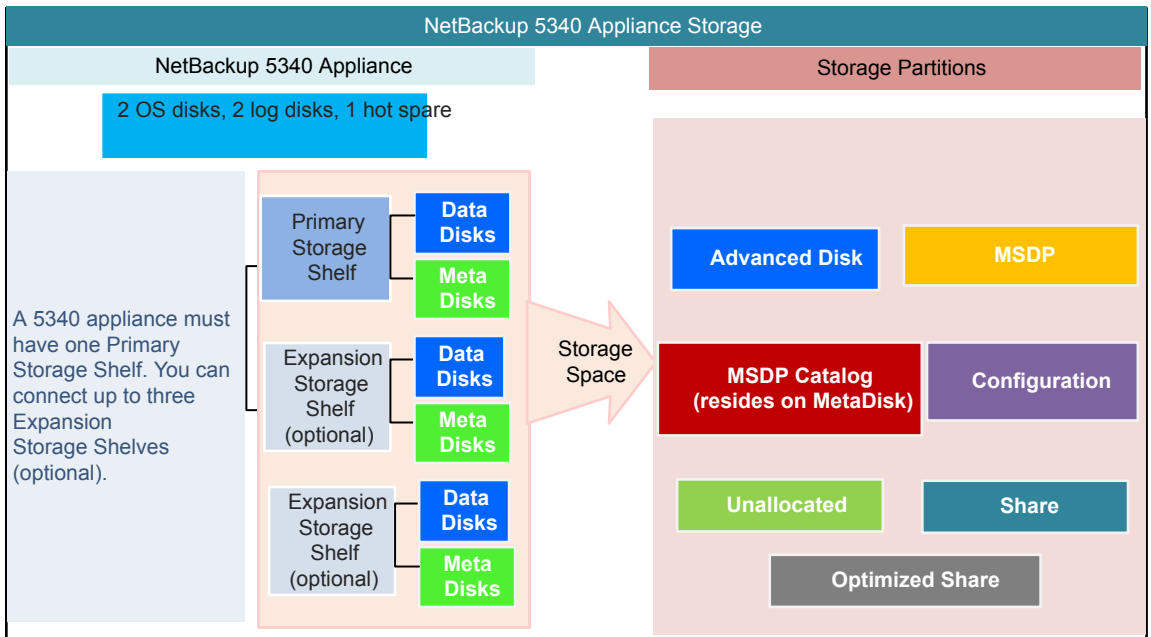


Figure 7-4 lists the tasks that you can perform on the appliance storage space.

Figure 7-4 Storage Operations

Storage Operations	
Tasks performed on Storage Disks	Tasks performed on Storage Partitions
<p>To perform the tasks listed below:</p> <ul style="list-style-type: none">- Go to Manage > Storage > Disks in the Appliance console.- Use the Manage > Storage shell menu	<p>To perform the tasks listed below:</p> <ul style="list-style-type: none">- Go to Manage > Storage > Partitions in the Appliance console.- Use the Manage > Storage shell menu
<p>Add</p> <p>Adds a disk in the New Available state. Adds disk space to the unallocated storage.</p> <p>Command - Add <Disk ID></p>	<p>Create</p> <p>Creates a share partition only.</p> <p>Command - Create Share <Standard/Optimized></p>
<p>Remove</p> <p>Removes disk space from the unallocated space.</p> <p>Command - Remove <Disk ID></p>	<p>Delete</p> <p>Deletes a share partition only.</p> <p>Command - Delete Share <ShareName></p>
<p>Scan</p> <p>Refreshes the storage disks and devices information.</p> <p>Command - Scan</p>	<p>Edit</p> <p>Edits the description and client details of a share.</p> <p>Command - Edit Share <Details> <ShareName></p>
<p>Show Disk</p> <p>Shows the disk's total and unallocated storage capacity and status.</p> <p>Command - Show Disk</p>	<p>Move</p> <p>Moves the partition from one disk to another.</p> <p>Command - Move <Partition> <SourceDisk> <TargetDisk> [Size] [Unit]</p>
<p>Tasks Common to Disks and partitions</p>	<p>Resize</p> <p>Create, resize, or delete a partition. You can delete a partition if Appliance is in a factory state (not configured as a master or media server).</p> <p>Command - Resize <Partition> <Size> <Unit></p>
<p>Monitor</p> <p>Displays progress of storage management tasks like Add, Remove, and so on.</p> <p>Command - Monitor</p>	<p>Show Partition</p> <p>Shows the partition's total, available, and used storage capacity. You can also view configuration and usage information for all partitions or specific partitions.</p> <p>Command - Show Partition <All/Configuration/Usage> [PartitionType] [Name]</p>
<p>Show Distribution</p> <p>Shows the distribution of partitions on a disk.</p> <p>Command - Show Distribution</p>	

All the tasks that can be performed on the NetBackup Appliance Web Console can also be performed by using the `Manage > Storage` shell menu.

For more information about `Main > Manage > Storage` commands, refer to *NetBackup™ Appliance Command Reference Guide*.

About configuring a shared storage pool

NetBackup Appliances are pre-tuned for optimum performance by Veritas. This section provides the best practices for architecting an environment with NetBackup 52xx series to obtain the most optimal performance. These best practices can help you avoid over-commitment during concurrent active stream counts, this is important to protect against I/O performance degrade in any storage pool. Lets consider the following examples to understand theses best practices being implemented in a common environment:

About the environment.

- There is a NetBackup domain consisting of two media servers and one shared storage pool. The storage pool uses independent storage units to direct backup operations from their associated clients, as well as to perform other operations such as duplications and image replication
- By default, the configuration of a storage pool resource does not limit the number of I/O streams (jobs) from all sources. This is due to the fact that storage pools of different types, technologies, and architecture can vary in performance.

Issues faced in the environment

- It is possible to over saturate a pool with I/O as a result of reading or writing too many streams from various operations including:
 - Backup jobs
 - Duplication jobs
 - Replication jobs
- This issue can occur regardless of the stream limitation settings for the individual storage unit configurations.
- Such over saturation issues tend to further worsen in case of SLP and AIR driven policies, due to the multiplicity of operations they can initiate to a single storage pool.
- When this over-commitment occurs, the performance of the target storage pool can degrade quickly, causing delays in operations, and potentially missed Service Level Agreements (SLA).

Recommendation 1:

The value for the **Maximum I/O streams** and **Maximum concurrent jobs** should be taken into consideration when calculating the total number of potential concurrent streams written to or read from a given storage pool from all operations such as backups, duplications, replications, and so on. For more information on the method to calculate the basic stream count for backups, See [“Calculating the basic stream count for backups”](#) on page 68.

In case of over saturation two key changes are recommended as best practices:

- Adjustment of the **Maximum I/O Streams** count configuration of the storage pool.
To configure the **Maximum I/O streams** value, select and edit the appropriate disk pool on the master server from the available list of disk pool devices in the NetBackup Administration Console.
 - On the NetBackup Administration Console follow the path **Media and Device Management > Devices > Disk Pools**.
 - At the bottom of the Change Disk Pool window, adjust the Maximum I/O Streams value by clicking on the check box, and specifying the value desired.
- Adjustment of the storage unit **Maximum concurrent jobs** value for each of the storage units that each media server uses to push data to the shared target storage pool.

To configure the storage unit's maximum concurrent jobs parameter, select and edit the appropriate storage unit on the master server from the available list of storage units in the NetBackup Administration Console.

- On the NetBackup Administration Console follow the path **NetBackup Management > Storage > Storage Units**.
- Adjust the Maximum concurrent jobs to a desired value based on the number of write operations for both backup and duplication jobs. The default value is set to '1'.

Note: The **Maximum I/O streams** value for the storage pool will take precedence if it is less than the combined storage units' 'maximum concurrent jobs' ability to send jobs to the shared storage pool.

Supposing that the two storage units sharing one disk pool have a combined max job count of 110 (with max concurrent jobs for each storage unit set to 55).

- If the **Maximum I/O streams** for the pool = 90, then the combined maximum job count allowed to be active will be 90.
- If 60 jobs are submitted to one storage unit and only 20 on the other, 5 jobs will still queue on the first storage unit even though the target storage pool could potentially handle them all

- Recommendation 2: An additional consideration can be made with respect to the design of the media servers and their associated storage units that point to the shared storage pool as follows:
- Currently the media servers are configured each with an independent storage unit, both of which push data to the same storage pool. This architectural setup can introduce a complexity which may multiply the number of write operations directed toward a single shared storage pool, especially if the storage pool's maximum I/O stream limit value has been defined
 - A simplification of the storage unit use can help assist in avoiding the complexity of the environment and potential over-commitment of the storage pool resource, especially when the source clients have many data selections or backup directives.
 - If the two media servers are reconfigured to share the same storage unit that points to the shared storage pool. This ensures the maximum number of concurrent write streams during backup and write-side of duplication operations are limited based by the storage unit's **Maximum concurrent jobs** value.

Calculating the basic stream count for backups

To calculate the basic stream count for backups multiply the following values together:

- Number of policies with concurrent backup windows
- Clients in each selected policy
- Backup selections in each selected policy
- Number of concurrent jobs per storage unit value
- Number of storage units using the given storage pool

Note: Additional stream counts should be considered when using storage lifecycle Policies to perform duplications, replications, and other related operations.

See [“About configuring a shared storage pool”](#) on page 66.

About moving a storage partition for better performance

This section describes how to move a storage partition from a base unit to an expansion disk. If all or a part of your MSDP partition resides on the Appliance base unit (base disk), it is recommended that you move the MSDP partition to an expansion disk. This applies to 52xx appliances and is needed for optimum performance. The base unit should ideally be used in case the external storage is full and you are yet to expand the storage to additional storage shelves.

While moving a storage partition from a base unit to an expansion disk the following guidelines should be considered:

- To schedule high performance backups, it is recommended to move the AdvancedDisk partition or MSDP partition from the base unit to the expansion disk . For example, moving the MSDP partition would result in the backup to MSDP to span across 450MB/s and 850MB/s storage which leads to performance improvements.
- The approximate time required to move 4TB from a base unit to an expansion disk is ~2.45hours on a light load system. This can be considered as the basis to calculate the time for moving additional TB of data.
- The amount of data being backed up should be limited when moving a storage partition. If the backup is not limited the time required to move a partition can increase and be double the actual time required. The best practice is to limit the backup coming in during the move, if this precaution is not taken the time to move 4TB the move can increase up to 5 hours

You can move a storage partition using the following options:

- Using the **Manage > Storage > Partitions** from the NetBackup Appliance Web Console, See [“Moving a partition”](#) on page 69.
- Using the `Manage > Storage > Move` command from the NetBackup Appliance Shell Menu, See [“Moving a partition using the NetBackup Appliance Shell Menu”](#) on page 71.

Moving a partition

This procedure describes the process to move a partition from one storage disk to another.

Note: The NetBackup Catalog partition cannot be moved. The NetBackup Catalog partition must always be present on the base unit of a 52xx appliance.

On a 53xx appliance, the MSDP Catalog partition must always be present on the Metadisk and can only be moved between Metadisks (wherever applicable).

To move a partition

- 1 Log on to the NetBackup Appliance Web Console.
- 2 Click **Manage > Storage**.
- 3 In the **Partitions** section, click the partition that you want to move. The partition details page opens.

4 In the **Partition Distributions on Disk** section, click **Move**.

See [“Move <partition> dialog”](#) on page 70.

5 Click **Move** to move the partition.

Note: The partition size and the workload on the system determine the time required to move a partition.

6 The Move dialog displays the progress details and status of the move operation.

Click **OK** once the operation is complete. The partition details page is automatically refreshed.

See [“About storage configuration”](#) on page 61.

Move <partition> dialog

The Move <Partition Name> window displays the following parameters:

Parameter	Description	Example
Source Disk	Displays the name of disk that currently holds the selected partition.	76YTG2BA7CBACB4F416D631CE (Base)
Partition Size	Displays the selected partition's size on the source disk.	300 GB
Target Disk	Click the drop-down list and select the target disk to which you want to move the partition.	9DB0FD2BA7CBACB4F416D631CE (Expansion)
Note: The Target disk must be different from the Source disk.		
Unallocated Size	Displays the unallocated size on the target device.	100 GB
Size	Type the storage size in GB, TB, or PB that you want to move from the current disk to the new disk.	35 GB
Note: It is an optional field. If the size is not specified, the appliance moves the entire partition.		
Note: The size to be moved cannot be greater than the Unallocated Size on the target disk.		

See [“Moving a partition”](#) on page 69.

Moving a partition using the NetBackup Appliance Shell Menu

In the NetBackup Appliance Shell Menu use the `Manage > Storage > Move` command to move storage partition from one disk to another disk. To move a partition from the NetBackup Appliance Shell Menu use the following procedure:

To move a partition from the NetBackup Appliance Shell Menu:

- 1 Log onto the NetBackup Appliance Shell Menu as an administrator.
- 2 To move a partition use the following command:

```
Move Partition SourceDiskID TargetDiskID [[Size] [Unit]]
```

In this command:

- *Partition* refers to the storage partition like AdvancedDisk, Configuration, MSDP and so on.
- *SourceDiskID* refers to the disk ID of the source disk. Use the `Storage > Show Distribution` command to view the disk ID and the partitions on the disk. For example, S0001FB362501ABDA0000009se.
- *TargetDiskID* refers to the disk ID of the target disk. Use `Storage > Show Distribution` command to view the disk ID and the partitions on it. For example, S2301FB362501ABDA0000009se.
- `[[Size]]` refers to the partition size. Partition size must be greater than 0. If you specify a partition size that is greater than the actual partition size, appliance moves the whole partition. The default value is [1].
- `[[Unit]]` refers to unit size (GB/TB/PB). The default value is [PB].

The following example displays the initiated procedure when you run the `Move` command:

```
Storage > Move AdvancedDisk S0001FB362501ABDA0000009se  
S2301FB3BC00540A62501ABDA0000009se 5 GB
```

```
- [Info] Performing sanity check on disks and partitions...  
(5 mins approx)
```

```
- [Warning] Do not reboot the appliance while the partition move is in  
progress.
```

Draft extracted 2013-9-20 10:06

```
Main > Manage > Storage view commands 111
```

```
Manage > Storage > Move
```

```
- [Info] The estimated time to move the partition can range from  
0 hours, 1 minutes to 0 hours, 2 minutes depending on the  
system load.
```

The greater the system load the longer it takes to complete the move operation.

```
>> Do you want to continue? (yes/no) yes
```

```
Moving part '1/1' disks... Done
```

```
- [Warning] No recipients are configured to receive software  
notifications.
```

Use 'Settings->Alerts->Email Software Add' command to configure the appropriate Email address.

```
- [Info] Succeeded.
```

Moving the MSDP partition from a base disk to an expansion disk for optimum performance

If all or a part of your Media Server Deduplication Pool (MSDP) partition resides on the appliance base unit (base disk), it is recommended that you move the MSDP partition to an expansion disk. This recommendation applies to 52xx appliances and is needed for optimum performance.

The following procedures explain how to move the MSDP partition from a base disk to an expansion disk. The base disk resides on the appliance base unit. The expansion disk resides on a storage shelf that is attached to the appliance. A 5220, 5230, or a 5240 appliance can have up to four expansion disks.

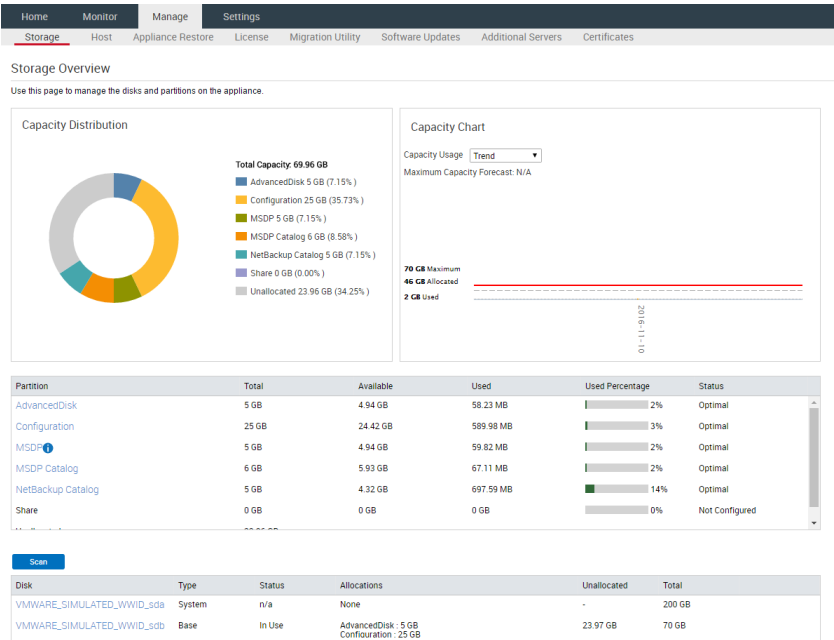
Consider the following scenarios:

- Scenario 1 - The MSDP and AdvancedDisk partitions are configured on the base disk. The expansion units are physically attached to the appliance but have not been added yet.
- Scenario 2 - The MSDP partition exists on the base disk. The expansion units are configured and partitions exist on them.

Select the scenario that applies to you and follow the appropriate procedure outlined below.

Scenario 1 - To move the MSDP partition from a base disk to an expansion disk

- Log on to the NetBackup Appliance Web Console.
- Click **Manage > Storage** and go to **Disks** section. Check the partitions that are on the Base disk. Suppose that you have MSDP, MSDP Catalog, AdvancedDisk, NetBackup Catalog, and Configuration partitions on the base disk.



- Ensure that the base disk is fully allocated by resizing the non-MSDP partitions (like AdvancedDisk). To ensure that the base disk is full, resize the AdvancedDisk partition to a value that is just below the maximum value.
 - In the **Partition** section, click **AdvancedDisk** partition to open the partition details page.

Partition	Total	Available	Used	Used Percentage	Status
AdvancedDisk	5 GB	4.94 GB	58.23 MB	2%	Optimal
Configuration	25 GB	24.42 GB	588.97 MB	3%	Optimal
MSDP	5 GB	4.94 GB	59.10 MB	2%	Optimal
MSDP Catalog	6 GB	5.93 GB	67.12 MB	2%	Optimal
NetBackup Catalog	5 GB	4.31 GB	708.43 MB	14%	Optimal
Share	0 GB	0 GB	0 GB	0%	Not Configured

- On the partition details page, click **Resize** in the **Partition Distributions on Disk** section. Enter a size in the **New Size** field that is slightly below the maximum value.

Veritas NetBackup Appliance

Connected to: Master 5220-174-eng128

Home
Monitor
Manage
Settings

AdvancedDisk

Capacity Distribution

Total Capacity

5 GB

Used Capacity

58.23 MB

Available Capacity

4.94 GB

Partition Distributions on Disk

Resize

Move

Partition Number	Disk	Type	Size
0	VMWARE_SIMULATED_VVWID_sdb	Base	5 GB

Resize Partition

The estimated time to resize the partition is 2 to 5 minutes.

Used Size:

58.23 MB

Unallocated Size:

23.0 GB

Current Size:

5 GB

New Size:

22.9

GB

Resize

Cancel

Click **Resize** to resize. In this example, the AdvancedDisk partition is being resized to 22.9 GB. The maximum value that it can be resized to is 23 GB.

Click **OK** after the resize operation is complete. The page is refreshed automatically and reflects the updated size.

- 4
- In the **Disks** section, click the **Add** link. Click **Yes** to confirm the addition and **OK** when it finishes. Repeat this process for the second expansion unit.

Storage Overview

Use this page to manage the disks and partitions on the appliance.

Capacity Distribution

Total Capacity: 69.96 GB

AdvancedDisk 22 GB (31.45%)

Configuration 25 GB (35.73%)

MSDP 5 GB (7.15%)

MSDP Catalog 6 GB (8.58%)

NetBackup Catalog 5 GB (7.15%)

Share 0 GB (0.00%)

Unallocated 6.96 GB (9.95%)

Capacity Chart

Capacity UsageTrend

Maximum Capacity Forecast: N/A

70 GB Maximum

46 GB Allocated

2 GB Used

2018-11-10

Partition	Total	Available	Used	Used Percentage	Status
AdvancedDisk	22 GB	21.81 GB	194.76 MB	1%	Optimal
Configuration	25 GB	24.42 GB	590.15 MB	3%	Optimal
MSDP	5 GB	4.94 GB	59.88 MB	2%	Optimal
MSDP Catalog	6 GB	5.93 GB	67.11 MB	2%	Optimal
NetBackup Catalog	5 GB	4.32 GB	696.51 MB	14%	Optimal
Share	0 GB	0 GB	0 GB	0%	Not Configured

Scan

Disk	Type	Status	Allocations	Unallocated	Total
VMWARE_SIMULATED_WWID_sda	System	n/a	None	-	200 GB
VMWARE_SIMULATED_WWID_sdb	Base	In Use	AdvancedDisk : 22 GB Configuration : 25 GB	6.97 GB	70 GB
VMWARE_SIMULATED_WWID_sdc	Expansion	New Available	None	-	200 GB
VMWARE_SIMULATED_WWID_sdd	Expansion	New Available	None	-	200 GB

AddAdd

Note that the Unallocated space increases.

- 5
- Check the space occupied by MSDP partition and how the MSDP partition is distributed across disks..

In the **Partition** section, click the **MSDP** link to open the MSDP partition detail page.

HomeMonitorManageSettings

StorageHostAppliance RestoreLicenseMigration UtilitySoftware UpdatesAdditional ServersCertificates

Storage Overview

Use this page to manage the disks and partitions on the appliance.

Capacity Distribution

Total Capacity: 469.9 GB

AdvancedDisk 22 GB (4.68%)

Configuration 25 GB (5.32%)

MSDP 5 GB (1.06%)

MSDP Catalog 6 GB (1.28%)

NetBackup Catalog 5 GB (1.06%)

Share 0 GB (0.00%)

Unallocated 406.9 GB (86.59%)

Capacity Chart

Capacity UsageTrend

Maximum Capacity Forecast: N/A

70 GB Maximum

46 GB Allocated

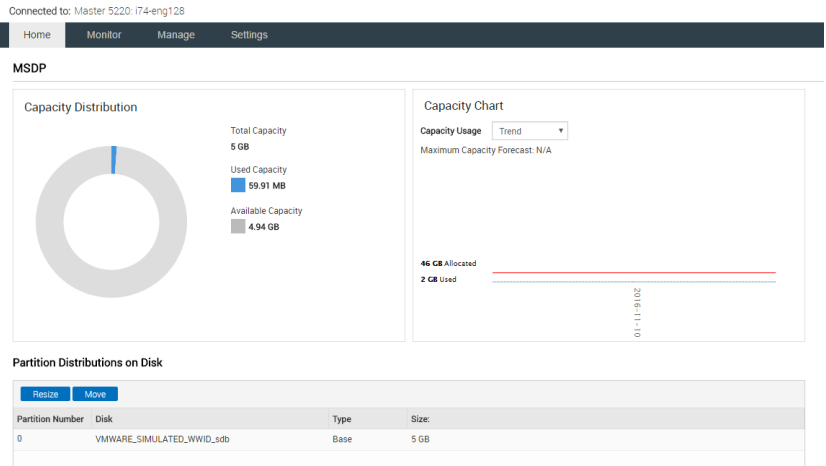
2 GB Used

Partition	Total	Available	Used	Used Percentage	Status
AdvancedDisk	22 GB	21.81 GB	194.76 MB	1%	Optimal
Configuration	25 GB	24.42 GB	590.41 MB	3%	Optimal
MSDP	5 GB	4.94 GB	59.91 MB	2%	Optimal
MSDP Catalog	6 GB	5.93 GB	67.11 MB	2%	Optimal
NetBackup Catalog	5 GB	4.32 GB	696.51 MB	14%	Optimal
Share	0 GB	0 GB	0 GB	0%	Not Configured

Scan

Disk	Type	Status	Allocations	Unallocated	Total	
VMWARE_SIMULATED_WWID_sdb	System	n/a	None	-	200 GB	
VMWARE_SIMULATED_WWID_sdb	Base	In Use	AdvancedDisk : 22 GB Configuration : 25 GB	6.97 GB	70 GB	
VMWARE_SIMULATED_WWID_sdc	Expansion	In Use	None	199.97 GB	200 GB	Remove
VMWARE_SIMULATED_WWID_sdd	Expansion	In Use	None	199.97 GB	200 GB	Remove

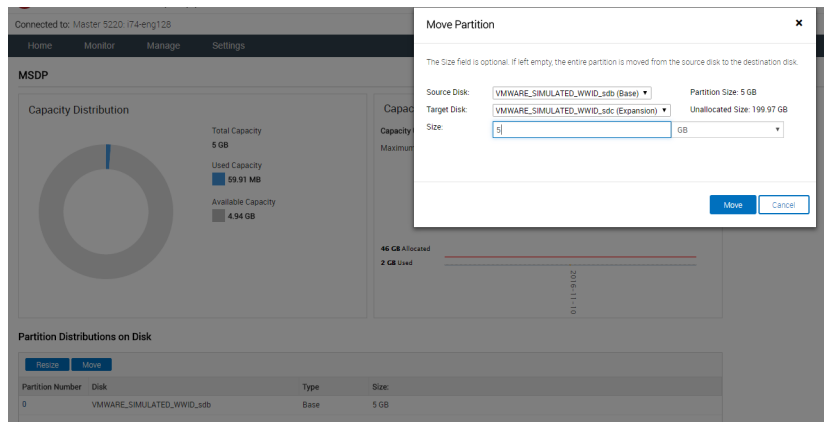
On MSDP partition detail page, check the **Partition Distributions on Disk** section.



In this example, all of the MSDP partition resides on the base disk and occupies 5 GB.

Note that the expansion disk must have at least 5 GB of unallocated space when you move the MSDP partition to the expansion disk at a later point.

- 6 Click **Move** in the **Partition Distributions on Disk** section.
- 7 The **Move Partition** window is displayed. Select the target disk from the drop-down list to which you want to move the partition.

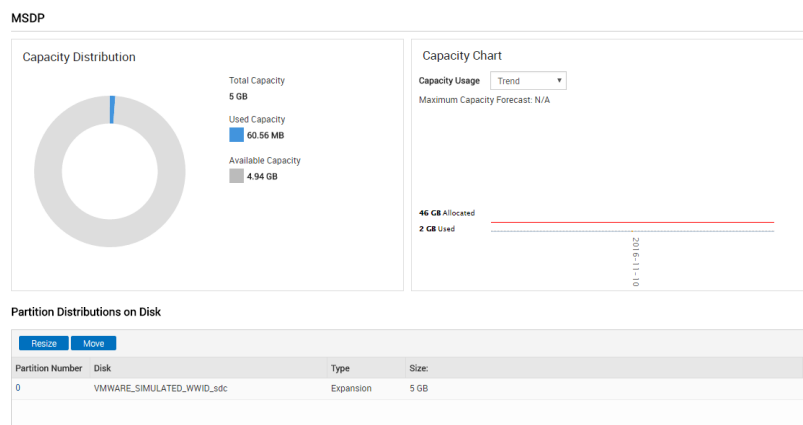


In this example, all of the MSDP partition that resides on the base disk is moved to one expansion disk.

- 8
- Click **Move** to move the partition.

Note: The partition size and the workload on the system determine the time it takes to move a partition.

- 9
- Dialogs in the **Move Partition** displays the progress details and status of the move operation. Click **OK** once the operation is complete. The MSDP detail page is automatically refreshed.
- 10
- On the MSDP detail page, check the **Partition Distributions on Disk** section. The MSDP partition resides on the expansion disk.

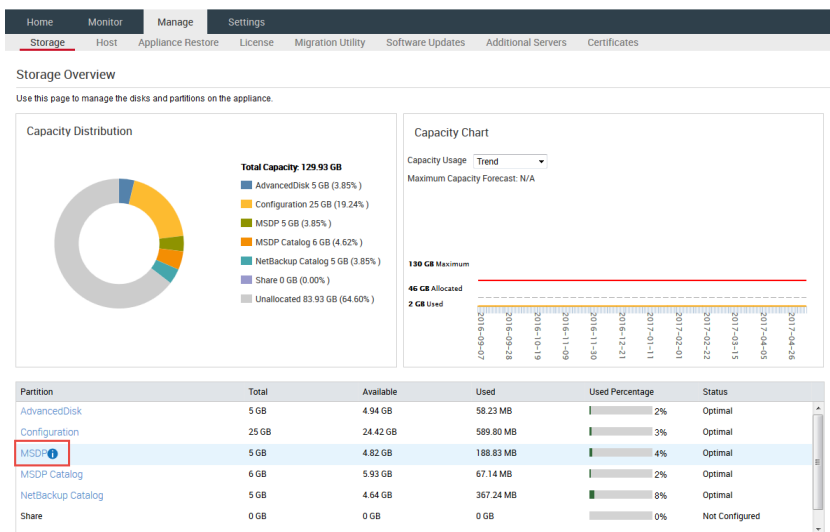


The following procedure explains how to move the MSDP partition from a base disk to an expansion disk when the expansion disk has partitions configured on it.

Scenario 2 - To move the MSDP partition from a base disk to an expansion disk

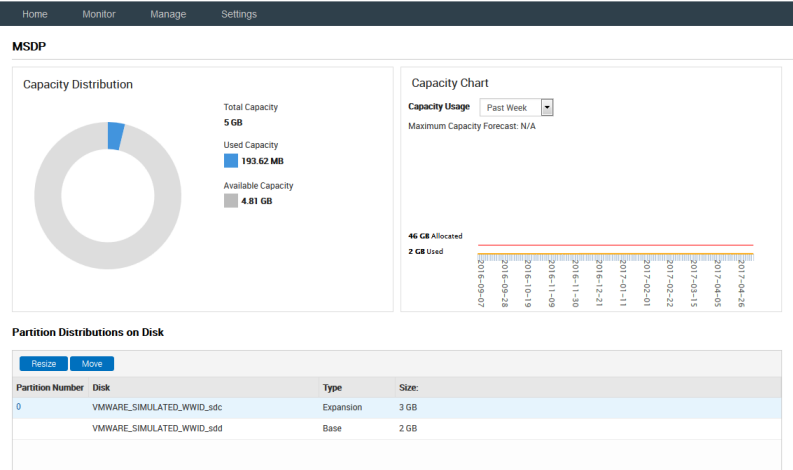
- 1 Log on to the NetBackup Appliance Web Console.
- 2 Click **Manage > Storage**.

In the **Partition** area, click **MSDP** to go to the MSDP detail page.



- On the MSDP partition detail page, check the **Partition Distributions on Disk** section.

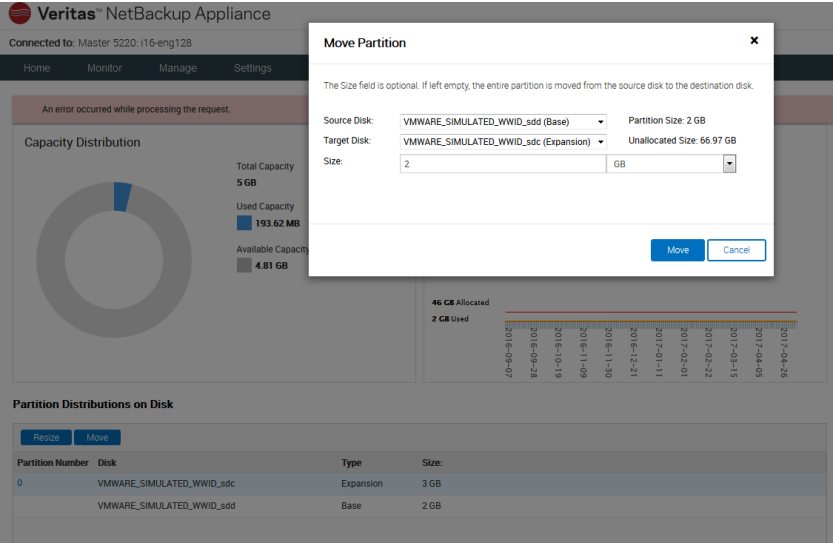
If the **Type** is Base for any of the disks, all or a part of the MSDP partition resides on the base disk. In this example, the MSDP partition is located on the base disk as well as the expansion disk.



If the **type** is Expansion for all the disks, the MSDP partition doesn't exist on the base disk. In this case, you do not need to move the MSDP partition. You can ignore the rest of the procedure.

- Click **Move** in the **Partition Distributions on Disk** section.

- The **Move Partition** window is displayed. Select the target disk from the drop-down list to which you want to move the partition.
- The target disk must be an expansion disk.



- Click **Move** to move the partition.

Note: The partition size and the workload on the system determine the time it takes to move a partition.

- After the operation completes, click **OK** to close the **Move Partition** window.

Generic best practices

This chapter includes the following topics:

- [Generic best practices](#)
- [About Notification settings](#)
- [About IPMI configuration](#)
- [Disaster recovery best practices](#)
- [Job performance](#)
- [Architecture](#)
- [NetBackup Catalog Backup](#)
- [Patching with the SLP \(storage lifecycle policies\)](#)
- [VMware backups using appliances](#)
- [Improving NetBackup Appliance restore performance](#)

Generic best practices

The generic best practices section includes a list of some good practices that can help you setup your appliance and improve its performance. This chapter includes the following sections:

- The About Notification Settings section describes the importance for enabling the notifications and re-directs you to register your appliance or edit the registration details on the [MyAppliance portal](#).
See “[About Notification settings](#)” on page 83.
- The IPMI section describes the importance of IPMI sub-systems and why it is vital to configure the IPMI system for your NetBackup appliance.

- See [“About IPMI configuration”](#) on page 84.
- The Disaster Recovery section provides the different configuration options that can help aid in a disaster recovery process if a disaster strikes.
See [“Disaster recovery best practices”](#) on page 85.
 - The Job performance section lists the best practices to improve the performance of the backup jobs run from your NetBackup appliance.
See [“Job performance ”](#) on page 86.
 - The Architecture section lists the best practices for defining your NetBackup Appliance architecture.
See [“Architecture”](#) on page 87.
 - The NetBackup Catalog Backup section lists the best practices for backing up the NetBackup Catalog from your NetBackup appliance.
See [“NetBackup Catalog Backup”](#) on page 87.
 - The Patching with the SLP section lists the procedure to perform patch updates when an SLP is involved.
See [“Patching with the SLP \(storage lifecycle policies\)”](#) on page 88.
 - The VMware backups using appliances section lists the best practices for performing VMware backups from your NetBackup appliance.
See [“VMware backups using appliances”](#) on page 89.
 - The Improving NetBackup Appliance restore performance section lists the two virtual memory tuning parameters to improve restore performance.
See [“Improving NetBackup Appliance restore performance”](#) on page 90.
- See [“About best practices”](#) on page 45.

About Notification settings

You can use the **Settings > Notification > Alert Configuration** from the NetBackup Appliance Web Console to apply the Call Home settings. AutoSupport in appliance uses the data that is gathered by Call Home to provide proactive monitoring for the appliance. If Call Home is enabled, the appliance uploads hardware and software information (or the Call Home data) to Veritas AutoSupport server periodically.

If the appliance encounters an error state, all hardware logs from past three days are gathered along with the current log. The logs are then uploaded to the Veritas AutoSupport server for further analysis and support. These error logs are also stored on the appliance. You can access these logs from `/log/upload/<date>` folder. If there is a problem with a piece of hardware, you might want to contact Veritas Technical Support. The Technical Support engineer uses the serial number of your appliance and assesses the hardware status from the Call Home data.

Note: For Call Home to work correctly, ensure that your appliance has Internet access either directly, or through a proxy server to reach the Veritas AutoSupport servers.

NetBackup appliance supports all the SNMP servers in the market. However, the following SNMP servers are tested and certified for using with version 3.1 and 3.1.1:

- ManageEngine™ SNMP server
- HP OpenView SNMP server

Also ensure that you register the appliance and your contact information on the [MyAppliance portal](#). Registering your NetBackup appliance helps to make sure that you are alerted to product updates and other important information about your appliance.

About IPMI configuration

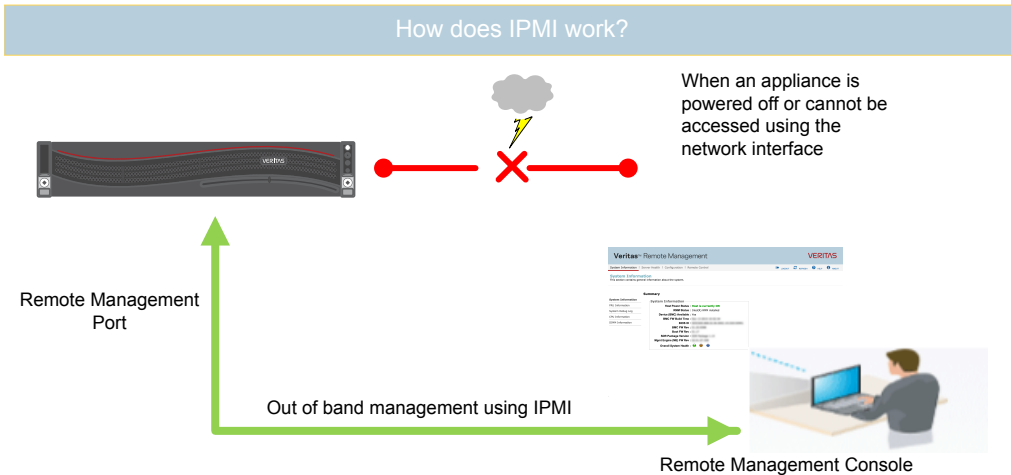
The Intelligent Platform Management Interface (or IPMI) provides management and monitoring capabilities independently of the host system's CPU, firmware, and operating system. You can configure the IPMI sub-system for your appliances. You can use the remote management port, located on the rear panel of the appliance, to connect to the IPMI sub-system.

The following figure shows the remote management port (or the IPMI port) on the rear panel of a NetBackup 5240 appliance:



The IPMI is beneficial after an unexpected power outage shuts down the connected system. In case the appliance is not accessible after the power is restored, you can use a laptop or desktop computer to access the appliance remotely by using a network connection to the hardware rather than to an operating system or login shell. This enables you to control and monitor the appliance even if it is powered down, unresponsive, or without any operating system.

The following diagram illustrates how IPMI works:



The following are some of the main uses of IPMI:

- Manage an appliance that is powered off or unresponsive. Using the IPMI, you can power on, power off, or restart the appliance from a remote location.
- Provide out-of-band management and help manage situations where local physical access to the appliance is not possible or preferred like branch offices and remote data centers.
- Access the NetBackup Appliance Shell Menu remotely using IPMI if regular network interface is not possible.

Note: Only the NetBackup Appliance Shell Menu can be accessed by using the IPMI interface. The NetBackup Appliance Web Console cannot be accessed by using the IPMI interface.

- Reimage the appliance from the IPMI interface by using ISO redirection.
- Monitor hardware health of the appliance from a remote location.
- Avoid messy cabling and hardware like keyboard, monitor, and mouse (KVM) solutions to access the appliance.

Disaster recovery best practices

NetBackup offers a few different configuration options that can help aid in a disaster recovery process if a disaster strikes.

Note: Use the following topology configurations as a general guide. Contact your Veritas account representative to establish what topology configuration best fits your particular environment.

Single domain configuration:

- Create backups of the MSDP catalog. The backup protects the critical MSDP information about the contents of the backup data that exists on the NetBackup appliance.
A policy is automatically created when configuring the NetBackup appliance for the first time as well as when adding MSDP storage during a Storage > Resize operation.
Review the policy configuration and make changes to its schedules, backup window, and residence as required. Make sure to activate the policy to protect the catalog.
See "MSDP catalog backup policy creation during initial configuration" in the *NetBackup Appliance 52xx Initial Configuration Guide* or the *NetBackup Appliance 53xx Initial Configuration Guide* for more information.
- Store catalog backups at an off-site location in case a recovery is necessary. You can use tape or cloud for restoration to a rebuilt master server at the disaster recovery site.

Multi-domain configuration:

- Configure Auto Image Replication to replicate backups that are generated in one NetBackup domain to storage in another NetBackup domain.

Job performance

This section provides the best practices for setting up backup and deduplication jobs.

Table 8-1 lists the best practices for running a backup or a deduplication job on your NetBackup appliance:

Table 8-1 Best practices to improve job performance

Best practice	Description
Distribute backup and deduplication jobs evenly:	<ul style="list-style-type: none">■ Distribute backup and deduplication jobs evenly across the available backup window.■ Do not front-load or back-load the job window.■ Limit the number of jobs and I/O at disk pool level for improved results in backup performance.

Table 8-1 Best practices to improve job performance (*continued*)

Best practice	Description
Monitor job activity	Monitor job activity and note any abnormal buffer times or delay. In case of any large abnormalities, contact Support for configuration changes that may potentially improve performance.
Evaluate Disk Polling Service (DPS) Proxy settings	Evaluate whether the Disk Polling Service (DPS) Proxy are set to the following: <ul style="list-style-type: none">■ <code>DPS_PROXYNOEXPIRE</code>■ <code>DPS_PROXY</code> (default <code>SEND/RECV</code> timeouts)
Avoid competing concurrent jobs like backup and restore (or Tape Out)	Competing concurrent jobs highly impact the job performance. Tape Out operations should be scheduled outside of backup window. Backup and replication can run concurrently up to the recommended 32 replication (Opt Dupe) streams without affecting performance significantly.

Architecture

This section lists the best practices to be considered for planning the network architecture around the Veritas NetBackup Appliance.

- Comprehensively gather the data protection requirements like SLA's, disaster recovery, retention periods, tape out requirements, and so on.
- Classify the data by type (for example, file system, database, images) to calculate proper and efficient deduplication rates.
- Segregate each data type to its own separate storage pool.
- Ensure 52xx Appliances do not exceed 85% capacity within the disk pool to ensure maximum performance.

See [“Generic best practices”](#) on page 82.

NetBackup Catalog Backup

The NetBackup catalog is the internal database that contains information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and the storage devices associated with its associated clients and the general infrastructure of the environment.

As NetBackup needs the catalog information so that it may perform restores of client backups, duplications, and other operations, it is important to configure a catalog backup policy before using NetBackup for normal day-to-day operations. It is also important to schedule the catalog backups to occur on a regular basis. Without regular catalog backups, an unexpected failure within the NetBackup may cause future operations to fail as well as a potential for a long and intensive recovery process by manual image import process.

The NetBackup appliance are designed to be a “drop-in”, disk-centric storage replacement and/or augmentation to existing traditional open-systems NetBackup architecture. As such, there are some additional considerations and best practices to ensure proper protection of the NetBackup catalogs in these environments:

- Configure the NetBackup Catalog backup policy to store the catalog images on a non-deduplication storage unit - ideally an AdvancedDisk pool.
- Always back up the catalog to a different storage host than the one local to the same master server. This helps to avoid losing the backup of the catalog, in the event of a system factory reset or system re-image.
- Back up the catalogs often to ensure subsequent backups, duplications, replications, and other catalog-related events are protected.
- Configure the Catalog backup policy to use email alerts to send the disaster recovery file to multiple mailboxes.

Note: For more information on NetBackup catalog refer to the *Protecting the NetBackup catalog* section from the *NetBackup Administrator's Guide for UNIX and Linux, Volume I* you can find the link to this guide on tech note [DOC5157](#).

See “[Generic best practices](#)” on page 82.

Patching with the SLP (storage lifecycle policies)

When planning to update a 52xx NetBackup appliance positioned as a media server in an existing NetBackup domain, it can be tricky to “halt” all operations directed to the appliance while the update procedure is occurring, especially when SLP operations are involved. An effective way is to suspend SLP operations by using the “suspend secondary operation” option from the NetBackup Administration Console on the master server for each operation directed to the intended 52xx media server.

To perform a patch updated, along with the SLP operations:

- 1** After the update is complete, deactivate the 52xx media server.
- 2** Suspend secondary operations for all SLP functions targeting the 52xx as a duplication destination.
- 3** Allow all the jobs to complete, or manually cancel them from the master server using the NetBackup Administration Console.
- 4** Proceed with the standard patching process as detailed in the patch release notes.
- 5** Reactivate the 52xx media server.
- 6** Reactivate SLP secondary operations.

See [“Generic best practices”](#) on page 82.

VMware backups using appliances

The following guidance is based on the example of an environment with a 10Gbps end to end bandwidth connection, a 60 hour window for full backups, and a 12 hour window for incremental backups. Using this scenario, additional solutions can be extrapolated based on scale:

- Based on average snapshot creation time frame of 1 minute (depending on the vCenter server, ESX & Datastore performance), the total number of snapshots supported per backup host (on an average) is 2500 (1 snapshot per datastore averaged simultaneously). The count of VM guests per proxy server shall be increased by load balancing jobs between ESX servers.
- The number of LAN-based media servers required (shared resources for virtual machines) to perform full backups of 2500 VM guests with 60 hour backup window will be approximately 2 ~ 4 considering the data to be backed up is roughly ~ 500TB. o In 8GB Fabric end-to-end environments, the number of media servers required will increase to 3 ~ 4.
- Considering a 5% change rate on a daily basis, the number of media servers required to perform an incremental backup of 2500 VM guests within a 12 hour backup window will be approximately. o In 8GB Fabric end-to-end environments, the number of guests can increase to approximately 2592.
- Considering an average of 20 VM guests per ESX host and 16 ESX hosts per cluster, the number of ESX clusters supported per backup host is approximately 8.
- The recommended number of snapshots per datastore is 1 (average 512 IOPS with 64K block size).

- The recommended number of backup requests per ESX server is 1.
- The average total number of jobs per cluster at any given point in time, based on the aforementioned recommendations is 48. However, based on best-practice recommendations, it is advised to limit the number of snapshots per vCenter server to 40 at any given point in time as it depends on the load factors of the vCenter server.

See [“Generic best practices”](#) on page 82.

Improving NetBackup Appliance restore performance

You can significantly improve the NetBackup Appliance restore performance if enough RAM storage is available to prefetch the fingerprint data.

To improve restore performance, you can change the following two virtual memory tuning parameters:

- Set `zone_reclaim_mode=0`.
Setting the `zone_reclaim_mode` parameter to zero may marginally degrade the backup performance in some cases because the memory allocation may take longer than in previous releases.
- Set `pagecache_limit_mb=0`.
Setting the `pagecache_limit_mb` parameter to zero can increase the swapping activities because it allows the file system page cache to grow. In this situation there is no limit on the size this cache can grow.

See [“About best practices”](#) on page 45.

Performance tuning

- [Chapter 9. Role-based Performance Monitoring](#)
- [Chapter 10. Optimize network settings and improve performance](#)
- [Chapter 11. Storage configurations](#)
- [Chapter 12. NetBackup Appliance tuning procedures and performance monitoring](#)

Role-based Performance Monitoring

This chapter includes the following topics:

- [Role-based configuration as a factor affecting performance](#)
- [What affects the performance of a Master Server?](#)
- [What affects the performance of a media server \(MSDP\)?](#)

Role-based configuration as a factor affecting performance

This section lists and describes the factors that how the role you configure your appliance as affect the performance of your appliance. NetBackup 52xx appliances can be deployed in the following backup roles in the datacenter:

- As a dedicated master server - See [“What affects the performance of a Master Server?”](#) on page 92.
- As a dedicated media server - See [“What affects the performance of a media server \(MSDP\)?”](#) on page 94.

Note: A NetBackup 53xx appliance is a media server by default and is not supported for the master server role configuration.

What affects the performance of a Master Server?

When an appliance is configured as a dedicated master server there are a number of parameters that can influence the performance and needs to be considered. lists

the parameters that need to be considered when you configure the roles as a master server.

Table 9-1 Parameters influencing the performance of a master server

Which are the process executed for this role?	Which are the resources used / affected due to the process?	How NetBackup 52xx help to optimally perform these processes?	Comments
The performance of a NetBackup master server depends on the following two diverse loads:	These processes require both CPU/Memory and disk I/O performance.	When you consider NetBackup appliance 52xx models, they provide the following features to optimize the resources used.	
EMM database	<p>For a master server to perform well, it depends on the performance of the EMM database, which is based on:</p> <ul style="list-style-type: none"> ■ Utilization of the complete CPU functionality and clock frequency The clock rate typically refers to the frequency at which a CPU is running. It is measured in the SI unit Hertz. ■ CPU and memory performance to handle the EMM database requests. 	<p>NetBackup 5220 has 2 CPUs with four cores and NetBackup 5230 has 2 CPUs with six cores.</p> <p>NetBackup 5240 has 2 CPUs with eight cores.</p>	Real world experience with NetBackup 5220 as a dedicated NetBackup Master server has demonstrated that appliance can handle 15000 backup jobs daily.

Table 9-1 Parameters influencing the performance of a master server
(continued)

Which are the process executed for this role?	Which are the resources used / affected due to the process?	How NetBackup 52xx help to optimally perform these processes?	Comments
File-based catalog information workloads	The performance of the disk I/O affects the master server's ability to execute and maintain the file-based catalog processes. The number of cores per CPU help to determine the type of disk I/O performance of a maser server.	<ul style="list-style-type: none"> NetBackup 52xx models have a formidable amount of RAM even in the basic configurations and very good disk I/O performance. This allows NetBackup 52xx appliances to play a role of the dedicated master server in a very busy environment. Scalability and diverse interfaces available in the appliances allow the unrestricted communication with the NetBackup Clients and the NetBackup Media servers and make NetBackup 52xx models are ideal candidates for master server from the connectivity performance perspective. 	

See [“Role-based configuration as a factor affecting performance”](#) on page 92.

See [“What affects the performance of a media server \(MSDP\)?”](#) on page 94.

What affects the performance of a media server (MSDP)?

NetBackup 52xx appliance as the dedicated MSDP media server is the most common type of deployment. In this kind of scenario there are many factors that influence performance of the media server:

There are two distinctive types of deduplication scenarios that can be major performance differentiators:

- **NetBackup Client deduplication**

In this scenario NetBackup client does part of the deduplication work. NetBackup Client will process backup data with the integrated deduplication plug-in (segments data and creates fingerprints) and send generated fingerprints to the NetBackup appliance for database lookup. This activity reduces the workload on the appliance and reduces the amount of data sent over the LAN from the

NetBackup Client to NetBackup Appliance while creating minimal load on the client.

- **NetBackup Media Server deduplication**

Media server-based deduplication places far greater workload on the appliance than the client-based deduplication since server processor is performing fingerprinting operations as well. Media server deduplication is a resource-intensive workload and running other concurrent operations that are resource-intensive will affect performance significantly. The appliance CPU cores play a primary role in the performance followed by RAM and disk I/O.

Table 9-2 Parameters influencing the performance of a media server

Which are the process executed for this role?	Which are the resources used / affected due to the process?	How NetBackup 52xx help to optimally perform these processes?	Comments
<p>In both types of deduplication scenarios the simplified workflow of standard deduplication load consists of the following tasks:</p> <ul style="list-style-type: none"> ■ Backup tasks ■ Segmenting ■ Hashing backup streams ■ Looking up in the deduplication segment Database – Metabase ■ Finally receiving and storing deduplicated segments on the disk 	<p>Deduplication as described places strain on the system resources</p> <ul style="list-style-type: none"> ■ CPU clock and CPU cores - Each backup stream is assigned to a single CPU core while the speed of processing of the streams is limited by CPU clock. Number of the CPU cores and sustained high CPU clock frequency are important for deduplication performance. ■ RAM - The RAM is important for processing of deduplication segments and for Metabase performance, ■ Disk I/O - while disk I/O have to be capable of handling incoming data processed by CPU and then sent to the file system where deduplicated segments are stored. ■ Veritas Filesystem (VxFS) - used for deduplicated segments storage on the appliance is capable of achieving high performance. It is important to keep in mind that file system performance will degrade as free space approaches 100% utilization, with 90% utilization being the recommended watermark level. 	<p>NetBackup 5220 Backup Appliance with its Intel Xeon quad core CPU and NetBackup 5230 with its latest generation Intel Xeon hexacore (six cores) CPU's provide powerful processing capabilities.</p> <p>NetBackup 5240 with its Intel Xeon CPUs with eight cores provide powerful processing capabilities.</p>	<p>To avoid resource congestion, ensure that when rehydration of data takes place no other concurrent operations run in parallel.</p> <p>During POC planning, the amount of client data has to be considered to not fill up the appliance file system and subsequently reduce both the backup and restore performance.</p> <p>All major subsystems of the appliance that are CPU, RAM, and Disk play an important role when it comes to deduplication performance. Running client-based deduplication which is lighter on the CPU than media server deduplication is a good practice, however careful consideration has to be taken to avoid running out of free memory.</p> <p>Note: In some specific cases there is a possibility that Client side deduplication performs slower than Media server deduplication. One such rare example is, backup of SQL Server transaction logs on the underpowered SQL Server. These cases are rare but possible and they are mentioned in the context of the exception that proves the rule.</p>

Table 9-2 Parameters influencing the performance of a media server
(continued)

Which are the process executed for this role?	Which are the resources used / affected due to the process?	How NetBackup 52xx help to optimally perform these processes?	Comments
Scheduled rebasing According to a predetermined schedule, once a day, rebasing is run to group segments from the same backup together on the disk to improve the restore performance.	Operation like rebasing has a very positive effect on performance, backup image segments are grouped together on the disk which improves disk I/O throughput.		It is important to keep in mind that file system performance will degrade as free space approaches 100 % utilization. The watermark level is 90% utilization.
Restore Restore operation requires full rehydration of data which is the most resource intensive process on the NetBackup MSDP.		NetBackup 52xx backup appliance can restore data using multiple restore streams.	
Tape out from MSDP	Tape Out operation is very resource-intensive operation and requires full rehydration of the backup image. Performance of the tape out operation is limited by how tape drives handle multiple streams. Unlike restore operations that can benefit from appliance performance with multiple concurrent restore streams, tape drives can handle only one stream per tape drive thus limiting tape out performance.		Limit the number of I/O streams to the recommended value of 96. This can be done in the properties of storage unit under Max I/O streams text box.

Table 9-2 Parameters influencing the performance of a media server
(continued)

Which are the process executed for this role?	Which are the resources used / affected due to the process?	How NetBackup 52xx help to optimally perform these processes?	Comments
CRQP (Content Router Queue Processing) It is one of the regular maintenance operations on MSDP that processes fingerprint database (CRDB) transaction logs (<code>tllogs</code>). This operation is necessary for the maintenance of the deduplication storage pool and is scheduled to run twice a day.	During CRQP activity there is a hit on the disk I/O and CPU/RAM resources in the range of 10% - 20% depending on the amount of tlogs to process and this has to be accounted for when planning for high performance tasks that will run concurrently with CRQP. Details of the CRQP activity can be monitored through spooled log: <code>/disk/log/spoold/storaged.log</code> . For more information, refer to TECH156743 for more details.		
Initial Optimized Duplication	Initial Optimized Duplication is an resource-intensive process and places a stress on the operating system. This severely affects the deduplication performance if executed concurrently with the backup.		

See [“Role-based configuration as a factor affecting performance”](#) on page 92.

See [“What affects the performance of a Master Server?”](#) on page 92.

Optimize network settings and improve performance

This chapter includes the following topics:

- [Optimize network settings and improve performance](#)
- [SAN Fiber Channel setup](#)
- [Network bonding](#)
- [VMware VADP](#)
- [Jumbo frame implementation for increased MTU](#)

Optimize network settings and improve performance

The Datacenter network environment is usually the most important factor when sizing for network performance. This section helps you to optimize the network settings for your appliance.

When considering the performance of an appliance you need to consider the following network parameters:

- SAN Fiber Channel setup - See [“SAN Fiber Channel setup”](#) on page 100.
- Network bonding - See [“Network bonding”](#) on page 101.
- VMware VADP - See [“VMware VADP”](#) on page 102.
- Jumbo frame implementation - See [“Jumbo frame implementation for increased MTU”](#) on page 103.

SAN Fiber Channel setup

NetBackup 52xx appliance can act as a powerful Fiber Transport media server ingesting data from SAN Client's using standard SAN Fiber Channel setup. From the performance perspective SAN Client provides one of the fastest methods to send data from NetBackup client to the appliance and presents a very resource-intensive workload on the appliance.

SAN Client can send data either to MSDP or AdvancedDisk on the appliance. Each of these destinations presents different workloads on appliance resources.

Table 10-1 Performance determinant with SAN Client setup

Implementation Type	Description	Guidelines
SAN Client with MSDP implementation	You cannot implement a NetBackup client-based deduplication when SAN Client is used.	<p>Media server deduplication mode is very resource-intensive on the appliance, especially on the CPU, and all guidelines for the Media server-based deduplication performance sizing apply in this scenario.</p> <p>In addition to the CPU workload caused by the deduplication process, the HBA itself presents additional load utilizing 5 Mhz of CPU clock cycles per every 1 MB/s of data throughput as a general rule of thumb.</p> <p>This load is significant and should be calculated as additional load on the appliance. Combined with Media server-based deduplication requirements, this will lower throughput calculation significantly based on how much bandwidth is used for the SAN Client data.</p>
SAN Client with AdvancedDisk implementation	AdvancedDisk workload primarily uses disk I/O resources of the appliance.	<p>The amount of data sent to the AdvancedDisk on the appliance should not be more than the appliance disk I/O capability.</p> <p>When sizing for the AdvancedDisk and SAN Client, maximum disk I/O throughput of 600 MB/s should be considered as a conservative number.</p>

Careful planning has to be done during architecture and performance sizing operations to achieve maximum performance. The following measures can be taken in addition to improve performance using a SAN Fibre Channel setup:

- Having additional HBA ports dedicated to FTMS will help with overcoming some limits on scaling (number of SAN Clients per port) and redundancy (some customers require redundant connectivity), but for the performance sizing limitation is still ingestion capability of the appliance.

- Best possible throughput of one SAN target port working in the 8 Gb SAN environment can reach 780 MB/s (1560 Mb/s for two FTMS ports in the target mode) when fully saturated. For 5220 this will represent maximum throughput possible on the PCI bus. NetBackup 5230 and 5240 appliance have a more advanced architecture and can handle more I/O bandwidth on the PCI bus.
- OS used with the appliances is based on the Linux OS and all FTMS or SAN Client limitations on Linux apply, for example number of streams to buffer size ratio.
- The number of FT streams on Linux and NetBackup Appliance is in direct relation with the buffer size and if the number of streams is increased the buffer size has to be reduced. Although it is possible to change the number of streams and reduce buffer size from the root prompt, deploying such a configuration in a production environment requires a support exception.

See [“Network bonding”](#) on page 101.

See [“VMware VADP”](#) on page 102.

See [“Jumbo frame implementation for increased MTU”](#) on page 103.

Network bonding

Network bonding has been available in the Linux environment for quite some time. This is a very useful network feature that significantly increases network scalability allowing multiple TCP streams to be load balanced across a number of network ports.

The commonly used network bonding modes with NetBackup Appliances are:

- **802.3ad bonding mode** - This mode works by balancing the TCP frames across switch ports.
- **balance-alb bonding mode** - This mode works by balancing the TCP frames across ports by using the operating system itself for frame balancing

They are both active-active in a sense that all NIC's or NIC ports are actively involved in the load balancing and provide load balancing for both incoming and outgoing data.

The following guidelines can help you to improve the performance of your appliance considering network bonding:

- Network bonding has to be done across interfaces of the same type, it is impossible to bond 1 Gb/s and 10 Gb/s NIC's together.
- Network Interface Cards do affect the CPU cycles to process data but this workload is not considered as significant on 1 Gb/s or even 10 Gb/s interfaces. The improved PCI communication and throughput in the NetBackup 52xx

appliance hardware has a positive affect on the CPU performance. These appliances allow for faster processing of the data bus interrupts with positive effects on CPU utilization during high network loads.

- When a number of 10 GbE NIC's bonded together and transferring data at the maximum will cause some CPU utilization. Exact CPU utilization is hard to quantify since it depends on a number of TCP parameters like MTU, network latency, and others.

See [“SAN Fiber Channel setup”](#) on page 100.

See [“VMware VADP”](#) on page 102.

See [“Jumbo frame implementation for increased MTU”](#) on page 103.

VMware VADP

The VADP (vStorage APIs for Data Protection) interface allows for direct connection to the storage array where virtual machines are stored. VADP enables backup products, like NetBackup, to perform a centralized, efficient, off-host LAN free backup of virtual machines. There are a number of variables in the VMware and NetBackup setup that can influence performance in this scenario. Some of these variables or considerations have been listed below

- VMware environment depends heavily on the storage performance and this is visible during backup operations.
- From the appliance side we should not oversee AdvancedDisk or deduplication capabilities of the appliance and enough leeway should be left for other workload resource utilization. For example HBA's and NIC's are using CPU resources as well and their usage has to be calculated into the overall workload.
- In case when MSDP storage on appliance is used to store backed up VMware data, the appliance performance will be limited by target deduplication processing capabilities of the appliance model and appropriate sizing should take this into consideration. All recommendations applicable to target deduplication apply to the VMware VADP backup to the MSDP pool on the appliance.

See [“SAN Fiber Channel setup”](#) on page 100.

See [“Network bonding”](#) on page 101.

See [“Jumbo frame implementation for increased MTU”](#) on page 103.

Jumbo frame implementation for increased MTU

Jumbo frames are the expression used to describe TCP frames with increased Maximum Transmission Unit – MTU to 9000. They can be very useful in improving throughput for some workloads and even reducing CPU utilization of the appliance.

Note: Jumbo frames are not covered by IEEE RFI's and are not considered network standard due to which they might not be supported by all network environments

Implementation of Jumbo frames ensures the following performance improvements:

- Increased MTU allows for faster transport of data when NetBackup client is sending large streams of data to the appliance, typically in the Media server based deduplication or the AdvancedDisk backup scenario. In these cases significant improvement in the LAN throughput is noticed.
- In order for the increased MTU to work, all other network devices in the same segment have to be configured with the same MTU, otherwise MTU will be negotiated down to the lowest MTU configured on the network device connected to the same network segment. This usually means that VLAN has to be created and that both NetBackup Client, all network equipment (switches, routers and other equipment) and the appliance have to be configured with the same MTU size.

Although increased MTU equates to a better throughput the following pointers need to be taken into consideration, prior to implementing Jumbo frames:

- NetBackup client based deduplication workload segments and hashes the data on the client and sends a number of smaller packets of data to the appliance. In this case, increased MTU size will not have a positive effect and should be avoided.
- Maximum MTU size that should be set is 9000 since this represents a limit of the TCP CRC (Cyclic Redundancy Check). A CRC is a method of checking for errors in data that has been transmitted on a communications link. Although this point is argued, it is without doubt that TCP CRC will have a risk of failing to protect data integrity if MTU is higher.

See [“SAN Fiber Channel setup”](#) on page 100.

See [“Network bonding”](#) on page 101.

See [“VMware VADP”](#) on page 102.

Storage configurations

This chapter includes the following topics:

- [Storage configurations](#)
- [Deduplication disk I/O and RAID level settings](#)
- [RAID Controller operations](#)
- [Deduplication load balancing](#)
- [Storage Lifecycle Policies](#)
- [Auto Image Replication \(AIR\)](#)
- [AdvancedDisk settings](#)
- [Tape Out operations](#)

Storage configurations

When considering the performance of an appliance you need to consider the following storage configurations:

- Deduplication disk I/O and RAID level settings - See [“Deduplication disk I/O and RAID level settings”](#) on page 105.
- RAID Controller operations - See [“RAID Controller operations”](#) on page 106.
- Deduplication load balancing - See [“Deduplication load balancing”](#) on page 109.
- Storage Lifecycle Policies - See [“Storage Lifecycle Policies”](#) on page 109.
- Auto Image Replication (AIR) - See [“Auto Image Replication \(AIR\)”](#) on page 109.
- AdvancedDisk settings - See [“AdvancedDisk settings”](#) on page 110.
- Tape Out operations - See [“Tape Out operations”](#) on page 110.

Deduplication disk I/O and RAID level settings

On the traditional storage array deduplication disk I/O and RAID level are bases for the performance measurement. With the NetBackup MSDP technology, as with vast majority of other deduplication technology currently on the market, deduplication disk I/O loses its primary performance measurement status. This being said, deduplication disk I/O still plays important part in the overall system performance, and should not be disregarded.

The following guidelines can help improve the deduplication disk I/O and RAID level settings:

- NetBackup 52xx appliance deduplication disk I/O comes under pressure during rehydration of the backup images stored in the MSDP or the AdvancedDisk backup/restore operations. These operations are heavily dependent on the deduplication disk I/O throughput.
- A traditional hard disk by its nature is very slow, mechanical devices cannot support concurrent I/O operations effectively. Storage array LUN's consisting of RAID groups are good in distributing load across a number of hard disks and thus improving disk storage I/O performance.
- RAID protocols have their own limitations and RAID 6 and RAID 5 have a disadvantage of slower performance during I/O write operations than RAID 1
- To further optimize the performance of the storage pool and avoid oversaturation, two key adjustments are recommended to the:
 - 'Maximum I/O Streams' count configuration of the storage pool. For the NetBackup 52xx appliance-based storage pools, Veritas recommends up to an optimal value of 98 streams per pool.
 - 'Maximum concurrent jobs' value for each of the storage units that each media server uses to push data to the shared target storage pool.

Note: Deduplication disk I/O throughput depends on the appliance storage configuration with the general rule of more disks meaning better performance. NetBackup 52xx appliance with 4TB storage configuration has a significantly lower deduplication disk I/O than a 24TB or better storage configurations. This is a general rule of thumb for any storage array and disk type, increased number of spindles provides better throughput at high disk I/O, especially when the number of streams is high.

See ["RAID Controller operations"](#) on page 106.

See ["Deduplication load balancing "](#) on page 109.

See ["Storage Lifecycle Policies"](#) on page 109.

See [“Auto Image Replication \(AIR\)”](#) on page 109.

See [“AdvancedDisk settings”](#) on page 110.

See [“Tape Out operations”](#) on page 110.

RAID Controller operations

RAID controller plays a very important role in the performance of the appliance. There is a number of operations performed by the RAID controller that can severely influence performance disk and they are analyzed in this section. [Table 11-1](#) lists the RAID processes performed by the RAID controller and the effect it has on the performance of the appliance:

Table 11-1 RAID Controller operations

RAID process	Description	Guidelines
RAID Rebuild Process	Disk storage on the NetBackup 52xx appliances is protected using RAID6 wherein a hot spare disk is made available on each shelf. In case of disk failure, the hot spare disk will be used to replace the data disk and a RAID rebuild process will be initialized.	<p>RAID rebuild is very demanding on the disk I/O and RAID controller resources. Since we are using 2TB and 3TB hard disks in appliances, there is a significant amount of data to move and a lot of RAID calculations to perform.</p> <p>This resource intense operation causes a major slowdown of appliance performance, affecting most appliance operations especially tape out, restore or AdvancedDisk operations, since they rely heavily on disk I/O for performance.</p> <p>Placing a load like Backup, Restore, Tape Out, etc. on the appliance while performing RAID rebuild will slow down the rebuild operation, and significantly extend the amount of time needed for the RAID rebuild to finish.</p>

Table 11-1 RAID Controller operations (*continued*)

RAID process	Description	Guidelines
BBU Learn Cycle Note: This process applies only to the NetBackup 5220 Backup Appliance.	NetBackup 5220 Backup Appliance has Li-ion Battery Backup Unit (BBU) used to protect the RAID controller write cache in the case of complete loss of power. If for any reason BBU is not operational, the RAID controller will switch off write cache (making it write-through) and this in turn will significantly slow down appliance disk I/O and affect performance of all operations on the appliance.	<p>BBU learn cycle usually takes around 3-4 hours to complete the full cycle of discharge and recharge. Appliance status displays that the battery is going through the learn cycle. Although this performance drop is unavoidable, it is a rare event since it happens once in 30 days and high performance operations should be planned around it.</p> <p>If RAID is going through the rebuild process it is recommended to reschedule BBU Learn cycle to avoid a slowdown in RAID rebuild performance. This can be done by the customer support engagement.</p> <p>It is possible to set amount of time until the next BBU Learn Cycle by contacting Support. The Support team can help you to set the number of seconds before next BBU Learn Cycle starts, counting from the last appliance restart. Although this method is not very exact, it allows for some kind of predictability and makes it possible to plan for high performance operations on appliance around this time.</p> <p>For more information on setting BBU Learn Cycle clock, refer to the tech note TECH204348.</p>
RAID consistency check	This is a regular consistency check operation performed on all LUN's on the RAID controller every 30 days. This check helps confirm that the integrity of RAID group is intact.	<p>It can influence the performance of the disk I/O operations but is not as significant as BBU Learn cycle. The default utilization of resources will be limited to 30% when the consistency check operation takes place.</p> <p>It is recommended to avoid overlapping high performance loads with consistency check. For example, if the BBU Learn cycle on the NetBackup 5220 Backup Appliance if overlapped with consistency check, it will slow down consistency check performance significantly. It is recommended to schedule these operations separately.</p>
RAID Patrol Read	RAID Patrol Read reads check physical hard disks for media errors. This operation is scheduled to run every seven days and is run simultaneously on a small number of physical disks at the same time.	<p>Patrol reads do not affect performance significantly since they are limited to the few disks at a time, but they still affect disk I/O on the checked disks.</p> <p>It is a best practice to make sure that patrol reads do not overlap other maintenance or recovery operations.</p>

RAID Controller commands

[Table 11-2](#) lists the commands to view the status of RAID processes:

Note: Ensure that all the MegaCLI commands are executed under the strict guidance of Support.

Table 11-2 RAID Controller commands

Description	Command
Monitor progress of the RAID rebuild	<code>/opt/MegaRAID/MegaCli/MegaCli64 -PDRbld -ShowProg -PhysDrv[Enclosure ID:Slot number of the hotspare disk] -aN</code>
Check whether consistency checker is running	<code>/opt/MegaRAID/MegaCli/MegaCli64 -ldcc -progdsply -Lall -aall</code>
To check whether Patrol Read is running:	<code>/opt/MegaRAID/MegaCli/MegaCli64 -adppr -info -aall</code>
To check BBU status:	<code>/opt/MegaRAID/MegaCli/MegaCli64 -adpBbuCmd -GetBbuStatus -aAll</code>
To configure RAID rebuild process to use maximum of RAID controller resources for the RAID rebuild (no operations should be active on the appliance):	<code>/opt/MegaRAID/MegaCli/MegaCli64 -AdpSetProp -RebuildRate 100 -a0</code>
After rebuild process finishes change should be returned to the default value):	<code>/opt/MegaRAID/MegaCli/MegaCli64 -AdpSetProp -RebuildRate 30 -a0</code>

See [“Deduplication disk I/O and RAID level settings”](#) on page 105.

See [“Deduplication load balancing”](#) on page 109.

See [“Storage Lifecycle Policies”](#) on page 109.

See [“Auto Image Replication \(AIR\)”](#) on page 109.

See [“AdvancedDisk settings”](#) on page 110.

See [“Tape Out operations”](#) on page 110.

Deduplication load balancing

Media server deduplication places a lot of pressure on resources and with implementation of Deduplication load balancing across multiple appliances or "build your own" media servers. This kind of load balancing can improve performance of deduplication processing on the appliance by offloading it to other servers but it is important to understand it's limits.

Although part of the CPU workload is offloaded to load balancing servers, memory utilization and disk I/O can become bottlenecks. Correct number of deduplication load balancing servers is usually reached by gradually incrementing number of additional servers until optimal performance is achieved. Two servers usually present optimal number but this might depend on the environment.

See ["Deduplication disk I/O and RAID level settings"](#) on page 105.

See ["Auto Image Replication \(AIR\)"](#) on page 109.

See ["AdvancedDisk settings"](#) on page 110.

See ["Tape Out operations"](#) on page 110.

See ["RAID Controller operations"](#) on page 106.

See ["Storage Lifecycle Policies"](#) on page 109.

Storage Lifecycle Policies

Storage Lifecycle Policies are widely used for duplication of backup images. Due to the nature of SLP this duplication happens 30 minutes after the backups have finished and many times this duplication process overlaps with other processes on appliance like backup, restore, tape out or other SLP process. SLP by its nature is not resource intensive once initial seeding is done. Initial seeding workload is very similar to restore workload, placing a lot of stress on the appliance resources and should be done over LAN instead of WAN. SLP performance can be significantly influenced by the high latency or high packet loss WAN connection.

See ["RAID Controller operations"](#) on page 106.

See ["Deduplication load balancing "](#) on page 109.

See ["Storage Lifecycle Policies"](#) on page 109.

Auto Image Replication (AIR)

NetBackup AIR is similar to SLP with the difference that NetBackup catalog is updated at the end of replication. This places additional workload on the Master

server on the destination side and in case this is all in one Master/Media server appliance this has to be accounted for during performance planning.

See [“Deduplication disk I/O and RAID level settings”](#) on page 105.

See [“RAID Controller operations”](#) on page 106.

See [“Deduplication load balancing ”](#) on page 109.

See [“Storage Lifecycle Policies”](#) on page 109.

See [“Auto Image Replication \(AIR\)”](#) on page 109.

See [“AdvancedDisk settings”](#) on page 110.

See [“Tape Out operations”](#) on page 110.

AdvancedDisk settings

AdvancedDisk is formatted across the appliance and storage shelves. It provides very good performance both for backup and tape out. It is used often if you have demanding tape out schedules and MSDP performance does not satisfy these strict requirements.

AdvancedDisk is formatted with VxFS and standard Storage Foundation. NetBackup tools can be used for the purpose of monitoring its performance (using commands like `vxstat`, `vxtrace`, `nbperfcheck` from the Netbackup Appliance Shell Menu).

See [“Deduplication disk I/O and RAID level settings”](#) on page 105.

See [“RAID Controller operations”](#) on page 106.

See [“Deduplication load balancing ”](#) on page 109.

See [“Storage Lifecycle Policies”](#) on page 109.

See [“Auto Image Replication \(AIR\)”](#) on page 109. See [“AdvancedDisk settings”](#) on page 110. See [“Tape Out operations”](#) on page 110.

Tape Out operations

Tape Out operation is very resource-intensive operation and requires full rehydration of the backup image.

Performance of the tape out operation is limited by how tape drives handle multiple streams. Unlike restore operations that can benefit from appliance performance with multiple concurrent restore streams, tape drives can handle only one stream per tape drive thus limiting tape out performance.

The following guidelines help to improve the performance of a Tape Out operation from MSDP:

- To achieve best results any concurrent operation is discouraged.
- Better performance improvements can be achieved if tape out is done after the rebasing operation is complete.
- It is very important to stream the data to tape drives on or above the tape drive minimum data rate. If stream throughput drops below the tape drives minimum data rate while tape out operation is in progress, the tape drive will stop writing and will start repositioning tape to the last saved segment while waiting for the new data. This is known as the “shoe shining” effect and will slow down the tape out performance significantly.

Note: Minimum streamed throughput of tape drives depends on the tape drive model and manufacturer. Generations of LTO 4 tape drives usually have minimum data rates of 40 MB/s (uncompressed data) and 80 MB/s (compressed data). Single MSDP rehydration streams have to perform better than this to prevent slowdown of the tape out performance. Another popular tape drive model is HP LTO5 drive (47 MB/s uncompressed and 94 MB/s for compressed data).

- Best possible scenario is to have multiple tape drives receiving multiple streams from appliance MSDP (one stream per drive).
- The number of tape drives to use depends on many environment parameters. If tape out restore speed is not satisfactory causing "shoe shining" on the tape drives, it is best to try to reduce the number of tape drives used out for tape gradually until an optimal number of tape drives is found for that particular customer environment.

See [“Deduplication disk I/O and RAID level settings”](#) on page 105.

See [“RAID Controller operations”](#) on page 106.

See [“Deduplication load balancing ”](#) on page 109.

See [“Storage Lifecycle Policies”](#) on page 109.

See [“Auto Image Replication \(AIR\)”](#) on page 109.

See [“AdvancedDisk settings”](#) on page 110.

See [“Tape Out operations”](#) on page 110.

NetBackup Appliance tuning procedures and performance monitoring

This chapter includes the following topics:

- [About diagnosing performance issues](#)
- [About performance tuning practices](#)
- [About tuning procedures and performance monitoring](#)

About diagnosing performance issues

When encountering a performance problem, the first step is to try reproducing the problem. If the problem is reproducible, the next step is to collect the four major system resource usage statistics while the problem is being reproduced. The four major system resources are:

- CPU
- Memory
- I/O
- Network

The most common utilities that are used to collect resource usage statistics are:

- **vmstat**
- **iostat**
- **ifstat**

With these three commands running concurrently, you can get a snapshot of the usage pattern of these four critical system resources. Examining the usage pattern of these resources could reveal the entry point to further performance issues. These utilities must be run concurrently with the same reporting interval. You can run these three utilities in the following manner.

Open three terminal windows and run each command in a separate window. In the NetBackup Appliance Shell Menu, go to **Support > Maintenance** menu to run these commands. Enter a password when prompted.

Note: These commands must be run concurrently with the same reporting interval.

Run the following commands:

```
vmstat 5 60 | tee /tmp/vmstat.5  
  
iostat -kxdt 5 60 | tee /tmp/iostat-kxdt.5  
  
ifstat 5 60 | tee /tmp/ifstat.5
```

When you run these commands, a copy of the data is saved in the `/tmp` directory.

Note: Ensure that there is enough space to save this data. Running out of space can disrupt normal operations.

Note that the number 5 after the command is the interval in seconds, while the second number 60 specifies how many reports to collect. You should not use the first report, because it gives the average statistics since the system was last started. You can change the interval (seconds) and the number of reports. In general, an interval of 5, 10, or 15 seconds is recommended for online monitoring. For a longer period of monitoring, you can use 30 seconds or 60 seconds to reduce the disk space needed to store the statistic files. Since the reports are the average over the interval, the best practice is to always collect the statistics using the same interval. This will help you develop a sense of how system resources are being consumed when the system is running normal. With this knowledge it will be easier to spot any change in resource usage and help find the root cause of a performance problem.

The following sections show how to detect potential performance bottlenecks from the output of the above three commands, and also the actions you can take to mitigate a performance problem. See the man page of these commands for detailed information about the output fields.

About CPU monitoring and tuning

Table 12-1 displays a sample output of the `vmstat` command when 120 streams of 98% deduplication backup jobs are running on a 53xx appliance.

Table 12-1 Sample `vmstat` output (collected with `vmstat 5`)

r	b	swpd	Free	Buff	Cache	si	so	us	sy	id	wa
89	0	1006344	34807856	37632	11694512	0	0	62	30	8	0
84	0	1006316	34845024	37640	12016276	11	0	62	30	8	0
63	0	1006316	34810400	37664	12260816	0	0	63	30	7	0
76	0	1006288	34785728	37664	12491148	5	0	61	29	9	0
46	0	1006288	34753840	37684	12756108	0	0	61	30	8	0
72	0	1006260	34711156	37692	13083760	3	0	62	30	8	0
72	0	1006252	34676620	37692	13332416	6	0	62	30	8	0
61	0	1006164	34648536	37712	13612680	28	0	59	29	13	0
92	0	1006156	34613540	37720	13902248	0	0	60	30	10	0
106	0	1006132	34572158	37724	14190992	6	0	61	31	9	0
82	0	1006128	34535448	37732	14465996	0	0	61	30	9	0
113	0	1005972	34507226	37740	14760008	30	0	61	30	10	0
66	0	1005964	34474724	37740	15004520	1	0	61	30	9	0
98	0	1005924	34444600	37748	15282376	8	0	60	30	10	0
118	0	1005920	34403518	37760	15582400	0	0	61	30	9	0
96	0	1005900	34380204	37764	15882380	4	0	62	30	9	0
60	0	1005900	34340626	37784	16175128	0	0	58	29	13	0
61	0	1005872	34303816	37792	16470724	3	0	62	30	7	0
60	0	1005868	34265376	37792	16747684	1	0	61	30	9	0
116	0	1005836	34234306	37800	17001952	5	0	62	30	8	0

Note: Some of the columns from the output have been removed to simplify the display.

From the above table, we can conclude that the system is CPU bound, because the **id** column (which displays the % of CPU idle) is mostly in single digit. This indicates that the 53xx CPU utilization is constantly over 90%. Another indication that the system is CPU bound is from the first column **r**. The value of column **r** is fluctuating between 46 and 118. **r** stands for "CPU ready to run queue". It is a count of processes that are currently running or ready to run but waiting for free CPU. 53xx has 40 logical CPU threads, it can at most handle 40 concurrent processes at a time. You can derive the number of processes that are ready to run but waiting for CPU cycles by subtracting 40 from the value in column **r**.

With the above CPU statistics and the fact that this happens while the system was running 120 concurrent 98% backup streams, there are two possible actions that you can take to lower the CPU consumption:

- Lower the batch size of job. If CPU is overly busy, the jobs could spend too much time waiting for available CPU cycles. Lowering the number of concurrent jobs per batch can improve overall performance.
- Adding another 53xx as the fingerprint server to double the CPU capacity is a natural solution.

A quick internal experiment with an additional fingerprint server showed that the performance increased almost 40% up to 10GB/sec while CPU usage reduced almost 50% on the appliance. At this point, the bottleneck switched to network since the 53xx can support up to 10 x 10 Gbps NIC which cap the network throughput around 10 GB/sec. We can probably see even higher performance improvement if there were more than 10 x 10 Gbps NIC installed on the system.

About memory monitoring and tuning

Besides monitoring CPU usage, `vmstat` output also displays the memory usage statistics. In the last example, we know memory is not the bottleneck because the **Free** column shows that there were over 340 GB of free memory. Also the **si** and **so** columns were near 0 mostly indicating that there's no memory swapping. These are a clear indication that there were no memory issue at the time of test. If **si** and **so** columns constantly show large non-zero values, for example, in the thousands, memory could be a possible bottleneck. You can use `top` command to check for processes that are consuming memory. If the problem is due to too many concurrent jobs running at the same time, try lowering the batch size to improve per job performance. Once a job is complete, the memory allocated for the process frees up which lowers the memory pressure and swapping activities.

About network monitoring and tuning

Table 12-2 shows the network statistics collected with **ifstat** when the system was running 120 concurrent 98% deduplication backup streams.

Table 12-2 Network Statistics (collected using ifstat 5)

	eth5		eth4		eth6		eth7		bond0	
HH:MM:SS	in	out	in	out	in	out	in	out	in	out
17:26:21	564730.9	1302.0	1160000.0	2140.2	559536.6	1621.4	168878.4	1153.4	4710000.0	12498.8
17:26:26	460546.3	728.0	1150000.0	698.9	459263.6	1330.3	381489.9	1142.0	4770000.0	11797.7
17:26:31	227735.0	537.6	1150000.0	724.0	445882.1	951.5	444105.8	1069.2	4680000.0	9971.4
17:26:36	374109.8	962.3	889411.5	863.5	421187.4	450.1	423359.6	1091.7	4740000.0	8593.0
17:26:41	413063.0	826.2	827849.0	765.3	415254.8	454.0	414780.8	758.3	4710000.0	6980.5
17:26:46	416389.0	802.9	822851.8	309.1	413860.3	425.2	409673.9	393.7	4690000.0	5694.1
17:26:51	413397.2	644.5	866869.2	587.3	343333.5	439.7	407649.9	343.9	4650000.0	5744.9
17:26:56	428205.4	510.6	1150000.0	2804.2	224508.3	952.0	427518.5	592.0	4550000.0	9509.6
17:27:01	474236.7	1040.3	1090000.0	2321.1	473665.8	1309.5	472919.9	832.9	4680000.0	12125.6

Note: Some of the columns from the output have been removed to simplify the display.

The above table showed the network throughput on the appliance of selected NIC ports. The appliance has 10 x 10 Gbps NICs bounded with the default balance-alb (adaptive load balance) mode. Under each NIC interface, the **in** column displays the kilobyte received per second from the network, while **out** column displays the kilobytes transmitted out per second. Eth4 is the primary slave which, in a couple 48 port switch models, tends to have a higher workload than the other slave NIC port. The data points highlighted in bold indicate that the NIC port is near capacity. With this condition, there is a possibility for packet loss. You can run the command `netstat -s |grep retransmit`, periodically to check for the number of retransmits. Nonzero value initially is fine, but if you see that this value increases, then there may be a network congestion issue.

The standard tuning recommendations are:

- Add or enable additional NIC ports

- Change the bonding mode, especially when all other ports, except the primary slave, still have plenty of extra bandwidth. The default balance-alb mode can sometimes cause this problem on some switch models. Our internal test showed 802.3ad (LACP) bond mode can solve this problem.

About I/O monitoring and tuning

[Table 12-3](#) is a sample output from **iostat**. Under the **Device** column, the **sdxx** displays the SCSI devices statistics, while **VxVMxx** display VxVM virtual volume statistics. See the man page of **iostat** command for a complete description of each column.

Table 12-3 Sample I/O statistics (collected with **iostat -kxdt 5**)

Device	Rrqm/s	Wrqm/s	r/s	w/s	rKB/s	wKB/s	Avgqsz	Avgqsz	Await	Svctm	%util
sdaw	0	0.4	0	8.8	0	3552	807.16	0.01	1.64	1.36	1.2
sdax	0	0	0	0	0	0	0	0	0	0	0
sdaz	0	0	0	0	0	0	0	0	0	0	0
sdba	0	0	0	8	0	2894	723.4	0.01	1.7	1.1	0.88
sdbb	0	0	0	0	0	0	0	0	0	0	0
sdbc	0	0	0	0	0	0	0	0	0	0	0
sdbd	0	0	0	0	0	0	0	0	0	0	0
sdbe	0	0.2	0	17	0	6786	798.33	0.03	1.88	1.36	2.32
sdbf	0	0	0	0	0	0	0	0	0	0	0
sdbg	0	0.4	0.2	14.8	1.6	5468	729.32	0.12	8.11	4.53	6.8
sdbh	0	0	0	0	0	0	0	0	0	0	0
sdbi	0	0.2	0.2	8.8	1.6	3222	716.27	0.02	2.67	1.69	1.52
sdbj	0	0	0	0	0	0	0	0	0	0	0
sdag	0	0	0	15.2	0	6358	836.63	0.03	2.05	1.32	2
VxVM3	0	0	0	0	0	0	0	0	0	0	0
VxVM4	0	0	0	0	0	0	0	0	0	0	0
VxVM5	0	45	0.6	17.8	4.8	12303	1337.85	0.04	2.26	1.61	2.96

Table 12-3 Sample I/O statistics (collected with `iostat -kxd 5`) (*continued*)

Device	Rrqm/s	Wrqm/s	r/s	w/s	rKB/s	wKB/s	Avgrqsz	Avquasz	Await	Svctm	%util
VxVM6	0	53.8	0	18.4	0	13502	1467.65	0.05	2.65	1.87	3.44
VxVM7	0	18	0.4	8.4	3.2	5743	1305.86	0.02	2.18	1.36	1.2
VxVM8	0	38.6	0.6	16.2	4.8	11225	1336.87	0.2	11.62	7.52	12.64
VxVM9	0	60	0.2	19.2	1.6	13064	1346.96	0.04	2.23	1.65	3.2
VxVM10	0	29.8	0	10.4	0	7349	1413.23	0.02	2.23	1.62	1.68
VxVM11	0	26.8	0.4	11.8	3.2	6573	1077.98	0.03	2.1	1.64	2
VxVM12	0	30	0.2	11.2	1.6	7440	1305.54	0.02	2.18	1.68	1.92
VxVM13	0	45	0.2	15.6	1.6	11652	1475.11	0.04	2.43	1.67	2.64
VxVM14	0	45	0.2	17.4	1.6	11895	1351.86	0.04	2.05	1.45	2.56
VxVM15	0	21	0.4	11.2	3.2	6814	1175.38	0.03	2.76	1.93	2.24
VxVM16	0	36	0.2	17	1.6	13358	1553.44	0.05	2.84	1.77	3.04

Note that `iostat` displays the I/O statistics for both SCSI and VxVM devices. Most of the time you only need to analyze the VxVM statistics. In the above table, VxVM5 – VxVM16 are the 12 VxVM volumes that MSDP used to store backup data. Each volume resides on a 7+2 Hardware RAID 6 LUN, and a file system is created on top of each LUN. So, each volume corresponds to one mounted file system.

The man page of `iostat` has a complete description of each column in the `iostat` output. The following describes some of the columns:

wrqm/s Write requests queued to the device that are merged per second. High number in this column indicates the I/O pattern is more sequential, and thus more opportunity for multiple I/O requests to be merged and send to the device as a single request. Most of the time high request merge for read and write improves I/O subsystem performance.

wKB/s Kilobytes written to the device per second.

avgrq-sz Average size (in sectors) of the I/O request

avgqu-sz Average number of I/O requests waiting on device queue

await Time in milliseconds for the I/O request to be served. This include both device queue time and the time the request was serviced (i.e. `svctm`).

svctm the average service time of I/O requests

wKB/s is KB written to the device per second and can be used to estimate the backup throughput per second. This can be done by adding wKB/s from VxVM5 – VxVM16 together. In a 0% deduplication workload, this sum should be very close to kilobytes received, i.e. “in” column under bond0 in the Network statistics table. Since the above statistics were captured while running 120 concurrent 98% deduplication backup streams, the wKB/s should be close to 2% of the kilobytes received from network. A simple calculation confirms the statement. The sum of wKB/sec from VxVM5 – VxVM16 is 12,097.1 KB, while in the Network statistics table, the value under bond0 “in” column is ~4,777,000KB. Dividing the two numbers (12,097.1/4,777,000) we get 2.5%. In other words, at this particular time, only 2.5% of backup data received from network need to be written to disks, because the rest of the data are duplicate of the existing data and there is no need to write them again.

In general, high deduplication ratio backup does not generate too much I/O activities, therefore it is unlikely to find I/O bottleneck. In the above table, disk service time, **svctm**, is mostly below 2ms, disk queue shows near 0 outstanding i/o, avgqu-sz, and %disk utilization are mostly below 5%. All these statistics indicate that I/O is not bottleneck. However, in the above table, we can see that for VxVM8 disk utilization %, svctm and await time are much higher than the other volumes. The statistics for this particular volume is not a cause for concern, you can safely ignore it. However, if you are curious about which file system the volume related to or if performance is bad enough that you want to find out more information about the volume, you can perform the following steps to find the filesystem that VxVM8 is mapped to:

1. Identify the device minor number by running the command `ls -l /dev/vx/rdisk/nbuapp`

The following is the sample output from the above command. The number in column 6, ie. “3 4 5 ... 12” are the device minor numbers, add VxVM in front of the number, you get VxVM3, VxVM4, .. VxVM12. These are the device name displayed in column 1 of the sample iostat output in [Table 12-3](#). The last column, 1pdvol, 2pdvol, ... 9pdvol are the VxVM virtual volume name which you can use for further drill down in step 2.

crw-----	1	root	root	199,	3	Sep	28	16:12	pdvol
crw-----	1	root	root	199,	4	Sep	28	16:12	1pdvol
crw-----	1	root	root	199,	5	Sep	28	16:12	2pdvol
crw-----	1	root	root	199,	6	Sep	28	16:12	3pdvol

```

crw----- 1      root    root    199,    7      Sep    28      16:12    4pdvol
crw----- 1      root    root    199,    8      Sep    28      16:12    5pdvol
crw----- 1      root    root    199,    9      Sep    28      16:12    6pdvol
crw----- 1      root    root    199,   10      Sep    28      16:12    7pdvol
crw----- 1      root    root    199,   11      Sep    28      16:12    8pdvol
crw----- 1      root    root    199,   12      Sep    28      16:12    9pdvol

```

- Identify the LUN used for the volume with the command `vxprint -ht`. The command prints the following for each of the pdvol. For example, the volume of interest in our case is VxVM8 which is corresponding to 5pdvol. The following output shows that 5pdvol resides on device nbu_appliance0_29. The output shows that 5pdvol resides on device 0_29. Prefix nbu_appliance in front of 0_29 to get the full device name (nbu_appliance0_29). This device name is needed in step 4 as a parameter to `vxdisk list` command.

```

v 5pdvol      -   ENABLED  ACTIVE   82039901904 SELECT -           fsgen
pl 5pdvol-01   5pdvol      ENABLED  ACTIVE   82039901904 CONCAT -           RW
sd 50002974280F356058357-01 5pdvol-01 50008000011F356058357 0 82039901904 0 appl 0_29 ENA

```

- Identify the file system that the volume, 5pdvol, is mounted on with the command “`df -hT`”. In the output of this command, you should find the following information corresponding to 5pdvol. The output shows 5pdvol is mounted on mount point /msdp/data/dp1/5pdvol and the size of the file system is 39TB.

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
/dev/vx/dsk/nbuapp/5pdvol	vxfs	39T	1.6G	38T	1%	/msdp/data/dp1/5pdvol

- Identify the SCSI device name with the command `vxdisk list nbu_appliance0_29`. At the end of output from the above command, you can find the following:

sds	state=enabled	type=secondary
sdav	state=enabled	type=primary
sddb	state=enabled	type=primary
sdbv	state=enabled	type=secondary

The above data shows that VxVM8 has four paths configured, two of them are active and two passive. In the `iostat` output, you should see that the `wKB/s` for VxVM8 is roughly the sum of `wKB/s` from `sdav` and `sddb`. In addition, the `wKB/s` from `sdav` and `sddb` should be very close to each other. This is due to the load balancing mechanism provided by VxVM multipath device driver, DMP.

General guidelines to spot a resource bottleneck

Identifying performance problem is not trivial, however, the following guideline can help quickly spot a potential problem:

- CPU may be a bottleneck if `us + sy` from `vmstat` output is consistently over 85% and/or the `r` column is more than four times the number of CPU cores on the system
- Memory may be a bottleneck if the value of “`si`” and “`so`” columns is consistently high in `vmstat` output (for example over 1000).
- I/O subsystem may be a bottleneck if the **await** column of VxVM device show greater than 25(ms). Another way to detect I/O bottleneck is continue increasing the load with an I/O intensive workload. When the throughput stop increases but the **await** time increases substantially, then it is very likely that the I/O subsystem is saturated.
- Network may be a bottleneck, if the sum of “`in`” + “`out`” is close to 1.2 GB/sec from `ifstat` output on a 10 Gbps or 120 MB/s on a 1 Gbps NIC.

For the first three items, consider lowering the batch size of concurrent jobs. Overall throughput may be better than overstressing the system because excessive jobs will most likely be on an internal queue somewhere waiting for resources to free up. For network bottlenecks, changing the bonding mode or adding more NIC ports should alleviate the problem.

About performance tuning practices

This section includes the following topics:

- [About I/O performance tuning](#)
- [About Oracle backup and restore performance tuning](#)
- [Setting NET_BUFFER_SZ to 0 on the NetBackup client for better WAN Optimization performance](#)

About I/O performance tuning

53xx has up to 12 file systems. In general, the best I/O performance can be achieved if I/O requests are well balanced across the 12 file systems. NetBackup provides the following three load balancing algorithms that you can choose from:

- 0 ---- disable load balance
- 1 ---- default load balance(by capacity and throughput)
- 2 ---- load balance by number of source ID

Default is 1 (load balance by capacity and throughput). The tuning parameter is set in `contentrouter.cfg` file under the directory

`/msdp/data/dp1/pdvol/etc/puredisk`. You can experiment with the other algorithms and find the best for your environment.

About Oracle backup and restore performance tuning

When you are backing up or restoring a large database on 53xx, the following tuning can substantially improve Oracle backup and restore performance:

1. Set **"ORACLE_OVERRIED_DATA_MOVEMENT = 1"**.

This setting disables the default proxy based mode when MSDP media is used. The parameter is set in the `bp.conf` file of NetBackup client under the directory `/usr/opensv/netbackup`.

```
NBU_Client# cd /usr/opensv/netbackup/
```

```
NBU_Client# more bp.conf
```

```
SERVER = nburdb01-bond0
SERVER = nbapp577-bond0
CLIENT_NAME = nburdb05-bond0
CONNECT_OPTIONS = localhost 1 0 2
ORACLE_OVERRIDE_DATA_MOVEMENT = 1
```

In stream-based mode, restore and backup can be processed in multiple parallel streams. With up to 12 mounted file system in 53xx and multiple RMAN channel configured, the Oracle backup/restore can improve substantially. In an internal test, stream-based mode outperforms proxy mode by over 300%.

2. **Optimal Database storage layout:** With the above tuning, 53xx Appliance can process substantial amount of Oracle data. However, to achieve the desired backup and restore performance, the I/O subsystem on the Oracle client must be able to perform at the same level. For example, to backup a 25TB database within 10 hours, the oracle client must be able to read at least 2.5TB of data

per hour and the network must be able to transmit at least 2.5TB of backup data.

3. **Choose appropriate RMAN setting:** Default RMAN backup sets `maxopenfiles=8`, which can create high I/O contention on the Oracle client. When `maxopenfiles=8` occurs in combination with 20 RMAN channels, there could be 160 parallel reads (8 * 20) on any given time during backup. To reduce the I/O contention, lowering the value of `maxopenfile` parameter can help improve overall performance. Internal tests with `maxopenfiles=1` yield better performance than the default.

Setting NET_BUFFER_SZ to 0 on the NetBackup client for better WAN Optimization performance

1. NetBackup 7.x by default set TCP receive buffer size to 32032 bytes on client node for restore job. TCP receive buffer size determines the amount of inbound data buffered by TCP stack while either waiting to properly sequence missing frames or waiting for the receiving application to read the already sequenced data. The following example is from tar log on RHEL6 platform, it shows the network receive buffer size used by the client node when running restore:

```
05:05:50 (259.001) Setting network receive buffer size to 32032 bytes
```

2. NetBackup appliance has provided an integrated WAN optimization component which aims to improve the outbound network transfer rate over WAN, its performance is highly related to the TCP receive buffer size according to the formula of "TCP throughput = TCP receive window size / Round-trip delay time". The default behavior of NetBackup client has posed a limitation to WAN Optimization performance with fixed TCP window size.
3. To achieve better WAN Optimization performance for restore job, placing a zero(0) into the `/usr/openv/netbackup/NET_BUFFER_SZ` file on the client node, which disables the TCP receive buffer tuning made by NetBackup, and lets the Operating System choose the proper value for TCP receive buffer size.

About tuning procedures and performance monitoring

Procedures and tuning parameters provided here are used in the field to improve the performance of certain appliance operations. Before tuning the appliance the following pointers should be considered:

- These parameters are not universally valid and they are not applicable for all deployments or scenarios. It is strongly recommended to test the appliance performance before and after tuning is done to make sure that there is an improvement.
- The provided commands in most cases do not have set parameters, rather they have a possible range or just a variable that needs to be replaced to fit a specific environment. It is highly recommended to understand and analyze commands and their effect before applying them.
- The provided output examples do not have real appliance performance data, they are given to provide information on the expected format of the output.

This section contains the following sub-sections:

- See [“NetBackup Client performance”](#) on page 124.
- See [“Other Performance Monitoring commands”](#) on page 127.

NetBackup Client performance

This section describes how to measure and improve the NetBackup Client side performance. There are a number of variables that can influence NetBackup client disk I/O performance. These variables need careful planning and approach when backing up high performance NetBackup clients. Some of the key variables that can influence NetBackup client disk are:

- Quality of storage array
- Non-backup related concurrent disk I/O workloads, like an antivirus scan running during backup cycle

Note: The disk I/O throughput on the NetBackup Client and NetBackup Appliance disk I/O throughput are different. The NetBackup Appliance disk I/O throughput is not related to the NetBackup Client performance.

- RAID groups hosting multiple LUN's with diverse workload and many others

As deduplication disk I/O represents a major choke point on the NetBackup Client side and it is very important to plan for an optimal number of backup job streams to saturate NetBackup Client I/O throughput. Rule of thumb is to use number of disks in the storage array LUN RAID group reduced by 2 to get the maximum number of streams. This rule is a good starting point but it is important to keep in mind that other datacenter environment variables will influence the outcome.

Using tools like `nbperfchk` to test NetBackup client disk I/O and troubleshoot potential bottlenecks on the data path to the appliance is essential during system tuning exercise.

Using nbperfchk to validate disk performance on NetBackup Appliances

The `Nbperfchk` command can also be used to test storage write speeds. If disk write speeds fall below 130 MB/s, there may be a potential issue with the underlying storage configuration in the appliance and Technical Support should be contacted to investigate. To test disk write speeds with `nbperfchk`, specify a path into the appliance storage being tested as the output path (`-o` argument value) for `nbperfchk`:

Table 12-4 Appliance storage paths

Appliance Storage	Storage Path	Example
52xx MSDP	/disk	<code>nbperfchk -I zero: -o /disk/nbperfchk.tst -s 64g -syncend</code>
52xx AdvancedDisk	/advanceddisk/advol	<code>nbperfchk -I zero: -o /advanceddisk/advol/nbperfchk.tst -s 64g -syncend</code>

Following are example runs of `nbperfchk` testing writes into MSDP and AdvancedDisk storage: 1) `Nbperfchk` write test into MSDP:

```
symmaster-a.Support> Nbperfchk run
Please enter options:
nbperfchk -I zero: -o /disk/nbperfchk.tst -s 64g
-syncend
Statistics log are recorded in nbperfchk_results.log
    3448 MB @ 1723.2 MB/sec, 3448 MB @ 1723.2 MB/sec, at 1369103472
    9380 MB @ 1875.6 MB/sec, 5932 MB @ 1977.3 MB/sec, at 1369103475
   12736 MB @ 1591.3 MB/sec, 3356 MB @ 1117.7 MB/sec, at 1369103478
   12736 MB @ 1157.1 MB/sec,    0 MB @    0.0 MB/sec, at 1369103481
   12736 MB @  909.1 MB/sec,    0 MB @    0.0 MB/sec, at 1369103484
   12736 MB @  748.6 MB/sec,    0 MB @    0.0 MB/sec, at 1369103487
   12736 MB @  636.3 MB/sec,    0 MB @    0.0 MB/sec, at 1369103490
   12736 MB @  553.4 MB/sec,    0 MB @    0.0 MB/sec, at 1369103493
   12736 MB @  489.5 MB/sec,    0 MB @    0.0 MB/sec, at 1369103496
   12736 MB @  438.9 MB/sec,    0 MB @    0.0 MB/sec, at 1369103499
   12736 MB @  397.8 MB/sec,    0 MB @    0.0 MB/sec, at 1369103502
```

```

12736 MB @ 363.7 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103505
12736 MB @ 335.0 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103508
12736 MB @ 310.4 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103511
12736 MB @ 289.3 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103514
12736 MB @ 270.8 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103517
12736 MB @ 254.5 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103520
18182 MB @ 342.8 MB/sec, 5446 MB @ 1815.4 MB/sec, at 1369103523
19936 MB @ 355.8 MB/sec, 1753 MB @ 583.9 MB/sec, at 1369103526
19936 MB @ 337.7 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103529
< output removed for brevity >
59456 MB @ 265.2 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103694
59456 MB @ 261.7 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103697
59456 MB @ 258.3 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103700
63330 MB @ 271.6 MB/sec, 3873 MB @ 1289.9 MB/sec, at 1369103703
Starting file sync to flush buffers...
Sync completed in 19.21 seconds
65536 MB @ 258.2 MB/sec

```

2) Nbperfchk write test into AdvancedDisk:

```

symmaster-a.Support> Nbperfchk run
Please enter options:
nbperfchk -I zero: -o /advanceddisk/advol/nbperfchk.tst -s 64g -syncend
Statistics log are recorded in nbperfchk_results.log
3776 MB @ 1886.3 MB/sec, 3776 MB @ 1886.3 MB/sec, at 1369103913
9104 MB @ 1819.1 MB/sec, 5328 MB @ 1774.4 MB/sec, at 1369103916
13036 MB @ 1627.9 MB/sec, 3931 MB @ 1309.2 MB/sec, at 1369103919
13036 MB @ 1183.9 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103922
13036 MB @ 930.2 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103925
13036 MB @ 766.0 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103928
13036 MB @ 651.1 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103931
13036 MB @ 566.2 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103934
13036 MB @ 500.9 MB/sec, 0 MB @ 0.0 MB/sec, at 1369103937
13529 MB @ 270.3 MB/sec, 493 MB @ 164.4 MB/sec, at 1369103961
< output removed for brevity >
59278 MB @ 315.0 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104099
59278 MB @ 310.1 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104102
59278 MB @ 305.3 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104105
59278 MB @ 300.6 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104108
59278 MB @ 296.1 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104111
59278 MB @ 291.7 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104114
59278 MB @ 287.5 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104117
59278 MB @ 283.4 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104120
59278 MB @ 279.4 MB/sec, 0 MB @ 0.0 MB/sec, at 1369104123

```

```

59278 MB @ 275.5 MB/sec,    0 MB @    0.0 MB/sec, at 1369104126
59278 MB @ 271.7 MB/sec,    0 MB @    0.0 MB/sec, at 1369104129
59614 MB @ 269.5 MB/sec,   336 MB @ 112.1 MB/sec, at 1369104132
63779 MB @ 284.5 MB/sec, 4164 MB @ 1387.7 MB/sec, at 1369104135
Starting file sync to flush buffers...
Sync completed in 18.65 seconds
65536 MB @ 268.4 MB/sec

```

All three examples show disk write throughput of more than 130 MB/s, which demonstrates good write performance for MSDP and AdvancedDisk storage.

Other Performance Monitoring commands

This section lists the other commands that can be used to monitor the performance of your NetBackup appliance.

- To check the VxFS (Veritas File System) fragmentation, use the following commands:

```

sudo /opt/VRTS/bin/fsadm -t vxfs -E <filesystem mount point>
sudo /opt/VRTS/bin/fsadm -t vxfs -D <filesystem mount point>

```

- To test disk I/O capabilities of the specific appliance using dd command, use the following commands:

Write performance testing:

```

nb5220:/home/maintenance # for i in 1 2 ; do time (dd if=/dev/zero
of=/cat/xyz$i bs=64k count=409600 oflag=direct); done

```

Read performance testing:

```

nb5220:/home/maintenance# for i in 1 2 ; do time (dd if=/cat/xyz$i
of=/dev/null bs=64k iflag=direct); done

```

See [“About monitoring the performance of your appliance”](#) on page 147.

See [“Role-based configuration as a factor affecting performance”](#) on page 92.

See [“Storage configurations”](#) on page 104.

Quick reference to Capacity planning and Performance tuning

- [Chapter 13. Capacity Planning checklist](#)
- [Chapter 14. Best practices checklist](#)
- [Chapter 15. How to monitor the performance of your appliance](#)

Capacity Planning checklist

This chapter includes the following topics:

- [Checklist for Capacity planning](#)
- [Capacity measuring worksheet](#)

Checklist for Capacity planning

What do you want to backup? - computer systems

The following questions provide the considerations to be made to backup computer systems.

Have you listed the Computer systems you need to backup? ☐

Have you noted the host name of each computer system to be backed up? ☐

Have you noted the operating system (and version) installed on each computer system? ☐

Have you noted the details about the Network technology used to connect the computer systems? ☐

Have you separately listed the computer systems connected to the network directly and the ones connected remotely? ☐

Have you noted the location of the resources to be backed up? ☐

Have you listed the total number of disks drives/tapes (and their model type) connected to each system? ☐

Have you listed the applications installed on the computer system? ☐

What do you want to backup? - Databases

The following questions provide the considerations to be made to backup databases.

Have you listed the Databases you need to backup? ☐

Have you noted the host name of each database? ☐

Have you noted the location of each database? ☐

Have you noted the operating system (version) installed on the database? ☐

Have you noted the details about the Network technology used to connect the database? ☐

Have you noted the database types and what are their version numbers? ☐

Have you noted the database engines and what are their version numbers? ☐

Have you noted the method used to backup these databases? ☐

Have you identified the databases that can be backed up while they are offline? ☐

Have you identified, for each database, whether the entire database or the raw partitions need to be backed up? ☐

What do you want to backup? - Application Servers

The following questions provide the considerations to be made to backup Application Servers.

Have you listed the Application servers you need to backup? ☐

Have you noted the host name of each Application server? ☐

Have you noted the server type and application release numbers of each Application server? ☐

Have you noted the location of each server? ☐

Have you noted the details about the Network technology used to connect the application server? ☐

Have you identified the application server that can be backed up while they are offline? ☐

Have you identified the method used to backup these application servers? ☐

How much data do you want to back up?

The following questions provide the considerations to be made for calculating the amount of data to be backed up:

Have you computed the total disk space on each individual system? ☐

Have you identified and added all the data types included in the backup? ☐

Have you identified how compressible is the data that you want to backup? ☐

Will the data be encrypted? ☐

Have you added the space on mirrored disks only once? ☐

Have you added the size of your database backups to your final calculations? ☐

Have you identified the data protection requirements? ☐

Have you added the requirements of the test-backup environment? ☐

Have you added any resulting growth factor to your total backup solution? ☐

Have you computed the amount of data you will need to backup in the next six months to a few years from now? ☐

When should the backup take place?

The following questions provide all the considerations for planning a backup schedule:

Will you perform daily-incremental backups? ☐

When will you perform a full backup? - Weekly ☐

When will you perform a full backup? - Bi-weekly ☐

When will you perform a full backup? - Monthly ☐

Have you identified the number of jobs to be run per day? ☐

Have you identified the window of time that is available to complete each backup? ☐

Would you want to configure a deduplication solution to take backups? ☐

Would you want to consider using the advanced NetBackup features such as synthetic backups, a local snapshot method, or FlashBackup? ☐

Do you plan to run an Off-site backup? ☐

Do you plan to run a Local backup? ☐

Have you identified the total number of local computer systems to be backed up? ☐

Have you identified the total number of remote computer systems to be backed up? ☐

Do you plan to backup any computer system over a network? ☐

Do you want to allow the users to run the backups and restores for their computer systems? ☐

Have you identified the policy for backup expiration? ☐

What is the retention period?

The following questions provide the considerations to be made when estimating the retention period:

What are the standards specific to your organization for creating a retention policy? ☐

Have you identified the projects requirement to retain the data? ☐

Have you identified the user-specific requirements on retaining data? ☐

Are their users who want to retain specific data for a longer period? ☐

Have you identified the business requirements specific to define the retention period? ☐

Have you considered the regulatory or the legal department laid requirements to retain the data? ☐

Have you calculated the size of your NetBackup image database? ☐

Designing your backup system - buying a new Media server appliance

The following questions provide the considerations you need to make for purchasing a NetBackup 52xx media server appliance:

Which is the Deduplication you want for the media server appliance? - File System ☐

Which is the Deduplication you want for the media server appliance? - VMware ☐

Which is the Deduplication you want for the media server appliance? - SQL ☐

Which is the Deduplication you want for the media server appliance? - Exchange ☐

Which is the Deduplication you want for the media server appliance? - Oracle ☐

What is the deduplication rates in percentage(%) applicable for Initial Full backups? ☐

What is the deduplication rates in percentage(%) applicable for Daily backups? ☐

What is the deduplication rates in percentage(%) applicable for Additional Full backups? ☐

How often are full backups taken? - Weekly ☐

How often are full backups taken? - Monthly ☐

Have you identified the number of incremental backups run per day? ☐

Have you estimated the data growth with each increment (%)? ☐

Have you estimated the Full backups to be taken per year? ☐

How long is the data retained? - Years ☐

Have you estimated the amount(MB) of Daily full backup to Advance Disk? ☐

Have you estimated the amount(MB) of Daily Incremental to AdvancedDisk? ☐

Have you identified the Retention Period (Days) for the AdvancedDisk? ☐

Have you identified the approximate yearly growth (%) for the AdvancedDisk usage? ☐

Have you estimated the Time Span (years) for the AdvancedDisk usage? ☐

For each deduplication use case, the following questions help to identify the Network parameters: ☐

Do you require SAN Transport connectivity for the media server? ☐

Do you required Fibre Channel connectivity for the media server? ☐

Do you require 2 Fibre Channel ports? ☐

Do you require 6 Fibre Channel ports? ☐

What is the estimated High ethernet throughput peak? ☐

What is the estimated Low ethernet throughput peak? ☐

Designing your backup system - buying a new master server appliance

The following questions provide the considerations you need to make for purchasing a NetBackup 52xx master server appliance:

Have you identified the total number of clients to be connected to the Master Server? ☐

How you identified the estimated data to be incrementally backed up in full (%)? ☐

Have you identified the total number of full backups the Master Server would take per Month? ☐

Have you identified the retention period for a full backup? ☐

Have you identified the total number of incremental backups the Master Server will take per Month? ☐

Have you estimated the approximate Total Number of files to be backed up? ☐

Have you identified the Average File Size (KB) to be backed up? ☐

Have you identified the total disk space used on all servers in the domain? ☐

Media and Master Server (All-in-one)

When you plan to buy a combined media and master server there are several considerations that you may need to make, so that you can identify the correct model that matches your requirements correctly. You will first need to identify all the parameters that the master server role of the appliance plays and then the parameters that the media server role of the appliance needs to satisfy.

Capacity measuring worksheet

Table 13-1 Capacity measuring worksheet

	File System	Exchange	SQL	Oracle	VMware
Initial Full (%)					
Full Dedupe (%)					
Incremental (%)					
TOTAL DATA (TB)					

Table 13-2 Capacity measuring worksheet

Information type	Information	
Amount of Data (TB)		
Number of Daily Incrementals		
Number of Weekly FULL		
Number of Daily FULL		
Monthly FULL		
Additional FULL		
Data Growth Annually (%)		
Retention Timescale Daily Incremental	Week	Month
Retention Period		

Table 13-2 Capacity measuring worksheet (*continued*)

Information type	Information	
Retention Timescale Weekly Full	Week	Month

Table 13-3 NetBackup appliance 53xx Performance sizing worksheet

[illegible]

Table 13-4 NetBackup appliance 5240 Performance sizing worksheet

[illegible]

Table 13-4 NetBackup appliance 5240 Performance sizing worksheet
(continued)

[illegible]**Table 13-5** NetBackup appliance 5230 Performance sizing worksheet[illegible]

Table 13-5 NetBackup appliance 5230 Performance sizing worksheet
(continued)

[illegible]**Table 13-6** NetBackup appliance 5220 Performance sizing worksheet[illegible]

Best practices checklist

This chapter includes the following topics:

- [Best practices checklist](#)

Best practices checklist

Table 14-1 Implementing deduplication solutions

Oracle	SQL	DB2	Sybase	Lotus notes
Agent				
<ul style="list-style-type: none"> ■ NetBackup for Oracle agent with RMAN (Stream-based backups) ■ NetBackup for Oracle Agent without snapshots (Proxy based backups) ■ NetBackup for Oracle Agent (Proxy based backups with Snapshot Client) ■ No NetBackup for Oracle Agent (Oracle Dump backups) 	NetBackup for SQL Agent	NetBackup for DB2 agent	NetBackup for Sybase agent	NetBackup for Lotus Notes agent
pd.conf settings for the media server				
	COMPRESSION = 1	COMPRESSION = 1	COMPRESSION = 1	COMPRESSION = 1

Table 14-1 Implementing deduplication solutions (*continued*)

Oracle	SQL	DB2	Sybase	Lotus notes
	MATCH_PDRO = 0			MATCH_PDRO = 0
		SEGKSIZE = 128	SEGKSIZE = 128	SEGKSIZE = 128
				LOCAL_SETTINGS = 0
Compression				
Turn off Oracle compression for the following agent: <ul style="list-style-type: none"> NetBackup for Oracle Agent with RMAN NetBackup for Oracle Agent without snapshots 		COMPRESSION turned off	COMPRESSION turned off	Disable: <ul style="list-style-type: none"> LZ1 compression Domino Attachment and Object Service Compress database design Compress document data
Defragmentation				
	Avoid frequent defragmentation. To be performed only when certain thresholds are reached.	Avoid frequent defragmentation.		
STRIPES				
Not applicable	Use multiple stripes.	Not applicable	Number of stripes has no impact on deduplication.	Not applicable
Other parameters				

Table 14-1 Implementing deduplication solutions (*continued*)

Oracle	SQL	DB2	Sybase	Lotus notes
<ul style="list-style-type: none"> For stream-based backups FILESPERSET = 1 in RMAN backup script for ASM-based backups For Oracle Dump backups leverage IMAGE MODE while dumping to local disk For stream-based backups and proxy based backups, increase number of channels. 	<ul style="list-style-type: none"> BLOCKSIZE = default BUFFERCOUNT = 2 buffers per stripe Leverage MAXTRANSFERSIZE Enable NO_COMPRESSION for SQL Server 2008 and higher 	<ul style="list-style-type: none"> Enable DEDUP_DEVICE parameter Data change rates <ul style="list-style-type: none"> Minimal data change rate, parallel streams = 1 High data change rate requiring high backup speeds <ul style="list-style-type: none"> Align number of containers in a DB2 database with the number of open sessions in the db2 database backup script 		Compaction - Run compaction at a scheduled interval weekly instead of user initiated compaction.

Table 14-2 Networking best practices

Consideration	Description
Generic best practices	<ul style="list-style-type: none"> Use Resilient Network feature of NetBackup to overcome remote office backup bandwidth issues related to latency and connection intermittence. Use 10GbE network infrastructure end-to-end wherever possible for maximum performance.
Fibre Channel implementations	<ul style="list-style-type: none"> Enable SAN Client - For high-speed backups and restores of NetBackup clients over Fibre Channel. Multipathing can provide SAN I/O high-availability and higher backup throughput. Enable Optimized Duplication to a NetBackup 5230 - To copy the backup images from a NetBackup 52xx appliance (source) to a NetBackup 5020 or 5030 appliance (destination). Enable NetBackup for VMware to backup and restore of the VMware virtual machines that run on VMware ESX servers. Implement Fibre Channel connection from NetBackup 52xx appliances to TLD tape storage devices allowing for backup data to be written to tape.

Table 14-2 Networking best practices (*continued*)

Consideration	Description
SAN zone configurations - FT Backup Zone	<ul style="list-style-type: none"> Only include the Fibre Transport traffic between the SAN clients and the appliance. Target port of the HBA—connect this port to a Fibre Channel switch port. In case of two target ports, define the appliance target port by physical port ID or World Wide Port Name (WWPN). Connect each SAN client HBA port to ports on the same Fibre Channel switch. Define the zones on the switch so that the client ports and the HBA target ports are in the same zone.
SAN zone configurations - FT Storage Zone	<ul style="list-style-type: none"> Carries the Fibre Transport traffic from a 52xx appliance to a 5020 deduplication appliance. initiator port of the HBA in the 5220 appliance—connect this port to a Fibre Channel switch port. Define the zones on the switch so that the 52xx appliance initiator port Only one initiator port and one target port can be configured for the same zone. Multiple initiator ports and target ports in the same zone are not supported.
SAN zone configurations - External tape storage zone	For a tape library as storage, create a separate zone for that traffic. The tape storage zone should include a port or ports on the FC HBA in slot 3 of a 52xx appliance.
Network bonding	<ul style="list-style-type: none"> Use network bonding for improved throughput and backup performance. It is a useful network feature that significantly increases network scalability allowing multiple TCP streams to be balanced across a number of network ports. The commonly used network bonding modes with NetBackup Appliances are: <ul style="list-style-type: none"> 802.3ad bonding mode - This mode works by balancing the TCP frames across switch ports. balance-alb bonding mode - This mode works by balancing the TCP frames across ports by using the operating system itself for frame balancing.
Network architecture	<ul style="list-style-type: none"> Comprehensively gather the data protection requirements. Classify and segregate the data by type. Ensure that 52xx Appliances do not exceed 85% capacity within the disk pool

Table 14-3 Working with storage partitions

Partition Name	Minimum supported size	Maximum supported size	Supported Platforms
AdvancedDisk	1 GB	Maximum available capacity	52xx 53xx

Table 14-3 Working with storage partitions (*continued*)

Partition Name	Minimum supported size	Maximum supported size	Supported Platforms
NetBackup Catalog	250 GB (Master server)	4 TB (Master server)	52xx (master server only)
NDMP Log	100 GB (Media server)	100 GB (Media server)	52xx 53xx
Configuration	100 GB	500 GB	52xx 53xx
MSDP	5 GB	Maximum available capacity	52xx
	5 GB	960 TB	53xx
MSDP Catalog	5 GB	25 TB	52xx, 53xx
Standard Share (Copilot)	5 GB	Maximum available capacity The limit for each individual share is 250 TB	5230, 5240, 53xx
Optimized Share (Copilot)	5 GB	114 TB or 228 TB	53xx
Universal Share	5 GB	Maximum available capacity	52xx 53xx

Note: To view the exact numbers for supported sizes, see the *NetBackup Appliance Product Description Guide* for the appropriate model.

Table 14-4 Generic best practices

Best practice	Description
Locate the NetBackup Appliance serial number	<p>Note and refer to the NetBackup Appliance serial number when you report an issue to support. You can use either of the following options to locate the NetBackup Appliance serial number.</p> <ul style="list-style-type: none"> ■ Monitor > Hardware > Health details from the NetBackup Appliance Web Console. ■ <code>Monitor > Hardware ShowHealth Appliance [Item] Command</code> from the shell menu.
Enable Call Home and AutoSupport	Use the Settings > Notification > Alert Configuration from the NetBackup Appliance Web Console to apply the Call Home settings. AutoSupport in appliance uses the data that is gathered by Call Home to provide proactive monitoring for the appliance.
Ensure that the IPMI sub-system is configured for your appliances	The Intelligent Platform Management Interface (IPMI) sub-system is beneficial when an unexpected power outage shuts down the connected system.
Enable Bare Metal Restore (BMR)	Use the Manage > Host > Advanced to enable BMR from NetBackup Appliance Web Console. BMR is the server recovery option of NetBackup that automates and streamlines the server recovery process. Thus making it unnecessary to manually reinstall the operating systems or configure hardware.
Perform Single domain disaster recovery configurations	<p>Create the appliance Master as a virtual machine, using Site Recovery Manager to replicate data from the protected site to the recovery site.</p> <p>Store catalog backups at an off-site location in case a recovery is necessary.</p> <p>Configure a setup using VCS Global Cluster Option for complete service level protection against an entire site failure by providing applications failover to the remote site</p>
Perform Multi-domain disaster recovery configurations	Configure Auto Image Replication to replicate backups that are generated in one NetBackup domain to storage in another NetBackup domain.
Improve Job performance	<ul style="list-style-type: none"> ■ Distribute backup and deduplication jobs evenly across the available backup window. ■ Limit the number of jobs and I/O at disk pool level for better performance. ■ Monitor job activity and note any abnormal buffer times or delay. ■ Evaluate the Disk Polling Service (DPS) Proxy settings - <code>DPS_PROXYNOEXPIRE</code> and <code>DPS_PROXY</code> (default SEND/RECV timeouts)

How to monitor the performance of your appliance

This chapter includes the following topics:

- [About monitoring the performance of your appliance](#)
- [Performance Monitoring matrix](#)

About monitoring the performance of your appliance

This section helps you to identify the processes and parameters that affect the performance of your NetBackup appliance. To help understand how to use your appliance with the most optimized settings, it is important to first understand the various factors that affect its performance. [Table 15-1](#) lists all the core factors that can help to improve the performance of your appliance, when configured in the correct manner:

Table 15-1 Factors affecting the performance of your NetBackup appliance

Factor Type	Factors	Description
Role configured	<p>When considering the performance of an appliance you need to consider the following roles:</p> <ul style="list-style-type: none"> ■ Master server - See “What affects the performance of a Master Server?” on page 92. ■ Media server - See “What affects the performance of a media server (MSDP)?” on page 94. ■ All-in-one Appliance - 	<p>The performance of your appliance greatly depends on the role it is configured as. For example, the performance of an appliance configured as all-in-one appliance will vary with the role of a media server or a master server.</p>
Network settings	<p>When considering the performance of an appliance you need to consider the following network parameters:</p> <ul style="list-style-type: none"> ■ SAN Fiber Channel setup - See “SAN Fiber Channel setup” on page 100. ■ Network bonding - See “Network bonding” on page 101. ■ VMware VADP - See “VMware VADP” on page 102. ■ Jumbo frame implementation - See “Jumbo frame implementation for increased MTU” on page 103. 	<p>The network architecture in which an appliance is used and the network settings applied to an appliance affect the backup and restore operations. For example, implementing a Fibre Transport media server to backup data from SAN Clients provides one of the fastest methods to backup data, however it also presents very resource-intensive workload on the appliance</p>
Storage configurations	<p>When considering the performance of an appliance you need to consider the following storage configurations:</p> <ul style="list-style-type: none"> ■ PureDisk I/O and RAID level settings - See “Deduplication disk I/O and RAID level settings” on page 105. ■ RAID Controller operations - See “RAID Controller operations” on page 106. ■ Deduplication load balancing - See “Deduplication load balancing” on page 109. ■ Storage lifecycle Policies - See “Storage Lifecycle Policies” on page 109. ■ Auto Image Replication (AIR) - See “Auto Image Replication (AIR)” on page 109. ■ AdvancedDisk settings - See “AdvancedDisk settings” on page 110. ■ Tape Out operations - See “Tape Out operations” on page 110. 	<p>The quantity of a data backed up and the rate at which data is backed up is greatly dictated by the type of storage you select that the corresponding storage configurations made for your appliance.</p>

Performance Monitoring matrix

Table 15-2 the following table lists the process that takes place in a Veritas NetBackup Appliance and its impact on the key resources of an appliance.

The relative impact on the resources is depicted using the following symbols:



Denotes that the resource (column) is highly impacted due to the corresponding (processes) listed in the rows.



Denotes that the resource (column) is moderately impacted due to the corresponding (processes) listed in the rows.



Denotes that the performance of the resource (column) is improved due to the corresponding (processes) listed in the rows.



Denotes that the resource (column) is not impacted or is not related to the corresponding (processes) listed in the rows.

Table 15-2 Process and resource impact matrix

Processes	CPU Cores	CPU clock	Memory	Disk I/O
Master Server				
Performance of the EMM database request				
File-based catalog information workloads				
Media server				
Backup tasks				
Segmenting				
Finger print operation of backup streams				
Looking up in the deduplication segment Database – Metabase				

Table 15-2 Process and resource impact matrix (*continued*)































































Processes	CPU Cores	CPU clock	Memory	Disk I/O
Receiving and storing deduplicated segments on the disk				
Rebasing				
Restore				
Tape out from MSDP				
CRQP (Content Router Queue Processing)				
Initial Optimized Duplication				
SAN Client				
MSDP implementation				
AdvancedDisk implementation				
Network Bonding				
1 GbE NIC's bonded together				
10 GbE NIC's bonded together				
VMware VADP				
Jumbo frame implementation				

Table 15-2 Process and resource impact matrix (continued)

Processes	CPU Cores	CPU clock	Memory	Disk I/O
RAID Rebuild Process				
BBU Learn Cycle				
RAID Consistency Check				
RAID Patrol Read				

See [“About monitoring the performance of your appliance”](#) on page 147.

Index

A

- about
 - master server role 39
 - media server role 40
 - NetBackup appliances 35
 - NetBackup documentation 9
- about the guide 8
- architecture
 - best practice 87

B

- back recording template 21
- capacity measuring 137
- best practice
 - architecture 87
 - catalog 87
 - checklist 141
 - generic 82
 - job performance 86
 - notification settings 83
 - SLP 88
 - VMware 89

C

- calculate basic stream count 68
- capacity measuring
 - back recording template 137
- capacity planning checklist 129
- catalog
 - best practice 87

D

- deduplication best practice 46
- DB2 backups 49
- generic recommendations 46
- Lotus Notes backups 51
- Microsoft SQL backups 48
- Oracle backups 47
- Sybase backups 50
- documentation 9

E

- enterprises use cases 31

H

- How much data to backup? 17

I

- intended audience 9
- IPMI configuration
 - about 84

J

- job performance
 - best practice 86

M

- maintenance sites 45
- master server
 - about role 39
- media server role 40
- move dialog
 - storage 70
- move partition
 - NetBackup Appliance Shell Menu 71
- move storage partition
 - improved performance 68

N

- NetBackup
 - about documentation for 9
- network considerations
 - fibre channel 53
 - introduction 52
 - SAN zone 53
 - validate network bandwidth 57

O

- optimize network 99
- jumbo frame 103

- optimize network (*continued*)
 - network bonding 101
 - SAN fibre channel 100
 - VMware VADP 102

P

- performance monitor
 - introduction 147
 - matrix 149
- performance monitoring
 - generic commands 127
 - NetBackup client 124

R

- remote or branch office use cases 34
- role
 - about master server 39
 - about media server 40
- role-based performance 92
 - master server 92
 - media server 94

S

- select master server 42
- select media server 41
- select new appliance 41
- shared storage pool 66
- SLP
 - best practice 88
- Solution introduction
 - High Availability 41
- storage configuration
 - about 61
- storage partition
 - moving 69

T

- template
 - application server information 27
 - computer system information 22
 - database information 24
- tuning procedures
 - introduction 123

V

- validate disk performance
 - nbperfchk 125

- VMware
 - best practice 89

W

- What is the retention period? 20
- What to backup? 14
- When to backup? 18