Symantec™ Disaster Recovery Orchestrator Deployment Guide

Amazon Web Services

6.1.1

November 2014

Symantec™ Disaster Recovery Orchestrator Deployment Guide

Product version: 6.1.1

Document version: 6.1.1 Rev 0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

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

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

www.symantec.com/business/support/index.jsp

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

Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.jsp

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

Contents

Chapter 1

Introduction to Disaster Recovery Orchestrator

This chapter includes the following topics:

Disaster Recovery Orchestrator overview

Disaster Recovery Orchestrator provides protection for the applications that are deployed in the IT setup of small and medium business (SMB) enterprises. The applications that are deployed on the on-premises site can be configured for monitoring and disaster recovery (DR). Such applications are migrated to or recovered in the Amazon cloud for which the SMB has a subscription.

A monitoring configuration protects an application against internal faults. If an application stops responding, the monitoring configuration attempts to restart the application and bring it online again. A DR configuration protects an application against site failures. If an application stops responding because the on-premises site becomes unavailable, the DR configuration can be used to recover the application in the cloud. An application that is configured with Disaster Recovery Orchestrator can also be migrated to the cloud.

Disaster Recovery Orchestrator components

Disaster Recovery Orchestrator components are classified as follows:

- Console components
  These components reside on a dedicated Amazon Elastic Cloud Compute (EC2) instance that acts as a controller for the disaster recovery (DR) activities.
- Client components
  These components reside on an on-premises system and its counterpart EC2 instance that manage the application monitoring activities.

Disaster Recovery Orchestrator Console

The Console components are installed on a dedicated EC2 instance in the cloud network, called the Console host.

The Console components and their functions are as follows:

- Disaster Recovery Orchestrator Authentication Service
  This service enables secure communication between the on-premises site and the cloud site. It uses digital certificates for authentication and SSL to encrypt communications. Disaster Recovery Orchestrator uses platform-based authentication; it does not store user passwords.
  The Console provides a single sign-on mechanism that uses this service so that an authenticated domain user does not have to:
  - Provide the user credentials associated with the EC2 instance to manage application monitoring.
  - Log on each time to connect to the EC2 instance.
- Role-based access control module
  This component manages the roles and users that are required for Disaster Recovery Orchestrator.
- Disaster Recovery Configuration wizard
  This wizard is used to configure an application for DR. When an on-premises application fails, you can continue servicing it from the cloud.
- Replication module
  This module manages the file replication mechanism that is used to synchronize the application data between the on-premises site and the cloud site.
- Symantec Storage Foundation Messaging Service (`xprtld`)
  The Console uses this service for the following functions:
  - Receive the application monitoring status, which is then displayed on the Dashboard and Application views.

- Relay commands that act on an application monitoring configuration, for example, takeover.

Disaster Recovery Orchestrator Client

Disaster Recovery Orchestrator Client is installed on the system where you wish to monitor an application. In your on-premises environment, the Client components can reside on a physical or a virtual machine. In the cloud, the EC2 instance on which you install the components is referred to as the cloud application host to differentiate it from the Console host.

The Client components and their functions are as follows:

- Application Monitoring Configuration wizard
  This wizard is used to configure application monitoring on the on-premises application host and on the corresponding cloud application host.
- Agents for configuring and monitoring applications
  - Infrastructure agents
    These agents manage resources such as heartbeats, NICs, storage, generic services, and so on.
  - Application-specific agents
    These agents manage the resources specific to those applications that are supported for monitoring, for example, Database Engine and FILESTREAM for SQL Server.
  - Replication agents
    These agents manage data replication between the on-premises site and the cloud site.
- Replication module
  This module ensures that application data at the on-premises site and the cloud site is in sync.
- Symantec Storage Foundation Messaging Service (`xprtld`)
  The infrastructure agents use this service to communicate the status of application monitoring to the Console.

Disaster Recovery Orchestrator licensing

Disaster Recovery Orchestrator follows a subscription-based licensing model. The licenses are metered on a per-instance basis, and the metering is done manually. An instance is defined as a 'protected application component'. If you change or renew the number of protected applications, you must report it to your Symantec Account Representative or your Symantec Certified Partner Reseller.

All licensing in Disaster Recovery Orchestrator is keyless. The Symantec product installer installs the embedded keys by default. A keyless license lets you use all the available product features.

For more information about the pricing, licensing, and the purchasing model, visit the Symantec website at:

https://licensing.symantec.com/

Deployment workflow

Deploying Disaster Recovery Orchestrator involves the following tasks:

1. Setting up the cloud infrastructure for disaster recovery (DR) readiness
   - Procuring an Amazon Web Services account
   - Creating Identity and Access Management (IAM) users and groups for managing the cloud resources
   - Creating user access keys for authentication
   - Creating an Amazon virtual private network (VPC) and defining its subnets and security groups for DR
   - (Optional) Creating an Amazon VPC and defining its subnets and security groups for performing fire drills
   - Launching an Amazon Elastic Cloud Compute (EC2) Windows instance to host the Disaster Recovery Orchestrator Console (Console host)
   - Launching Amazon EC2 Windows instances to host the Disaster Recovery Orchestrator Client and the applications to be configured for DR (cloud application hosts)
2. Installing the Disaster Recovery Orchestrator components
   - Allocating users and groups for installing and configuring Disaster Recovery Orchestrator
   - Installing Disaster Recovery Orchestrator Console on the Console host
   - Installing Disaster Recovery Orchestrator Client on the on-premises application hosts
   - Installing Disaster Recovery Orchestrator Client on the cloud application hosts
3. Configuring users and security settings
   - Creating users and assigning them privileges on the on-premises application hosts for the application configurations

- Creating users and assigning them privileges on the cloud application hosts to configure applications for monitoring and for recovery
  For more information, see the *Symantec Disaster Recovery Orchestrator Administration Guide*.

Common terms used in the context of Disaster Recovery Orchestrator

The Disaster Recovery Orchestrator solution caters to on-premises and cloud environments, and therefore, deals with a wide range of entities. Some of these entities may be referred to using multiple names. This section describes some common terms and conventions that are used throughout the Disaster Recovery Orchestrator user interface and documentation. These terms are listed in an alphabetical order.

Amazon Web Services (AWS)

Disaster Recovery Orchestrator uses the following AWS entities:

- Amazon Elastic Block Store (Amazon EBS) volume
- Amazon Elastic Cloud Compute (Amazon EC2) instance
- Amazon Machine Image (AMI)
- AWS Identity and Access Management (IAM)
- AWS Management Console
- availability zone
- customer gateway
- paravirtual (PV) virtualization
- security group
- virtual private cloud (VPC)
- virtual private gateway (VPG)

For information about these entities, see the AWS glossary:

http://docs.aws.amazon.com/general/latest/gr/glos-chap.html

application monitoring

A feature of Disaster Recovery Orchestrator that enables you to monitor applications running on a physical computer or a virtual machine. If the application components

fail and cannot be recovered after a certain number of attempts, the application status is reported accordingly on the Health View or the Console UI.

application monitoring configuration

An application that is configured for monitoring on the on-premises application host or the cloud application host. Disaster Recovery Orchestrator Client manages the application monitoring configuration on a system.

application recovery configuration

An application that is configured for migration or for disaster recovery (DR) in the cloud in the event of a failure at the on-premises site. Disaster Recovery Orchestrator Console manages all the application recovery configurations. An application must be configured for monitoring before it can be configured for DR.

The following systems are associated with every application recovery configuration:

- On-premises application host
- Console host (a single Console instance caters to multiple application recovery configurations)
- Cloud application host

These systems are also referred to as the systems that participate in the DR solution for an application.

cloud application host

The Amazon EC2 instance on which Disaster Recovery Orchestrator Client is installed. This instance acts as the in-cloud counterpart of the on-premises system that hosts the application that is configured for recovery. The application is configured for monitoring on this instance as well. If the on-premises application or its host becomes unavailable a recovery administrator performs Takeover. When Takeover is successful, this instance begins running the application.

Client

The Disaster Recovery Orchestrator client component that manages the authentication, file replication, application monitoring, and user interface modules.

Client host

The on-premises system or the Amazon EC2 instance on which Disaster Recovery Orchestrator Client is installed. This system hosts the application and its monitoring configuration.

Console

The Disaster Recovery Orchestrator server component that manages the authentication, file replication, DR, and user interface modules.

Console host

The Amazon EC2 instance on which Disaster Recovery Orchestrator Console is installed. This instance acts as controller for the DR activities.

Console UI

The Disaster Recovery Orchestrator Console user interface, which is browser-based.

failback

The operation in which application processing is restored on the original on-premises system when it becomes available again. A recovery administrator manually triggers Failback, but the tasks that are involved in the operation are performed automatically in a predefined sequence.

file replication

The replication mechanism that Disaster Recovery Orchestrator uses to synchronize application data between the on-premises site and the cloud.

fire drill

A feature of Disaster Recovery Orchestrator that lets you test your DR configuration. A fire drill operation tests the Takeover operations on an application that is configured for DR. When a fire drill is successful, the application comes online in a separate virtual network in the cloud, without disrupting the application in the production environment.

guest user (role)

A user who has the privileges to view the application recovery configuration on a system. A security administrator adds this role for a user to the Privilege Settings tab of the Console UI.

journal file

The intermediate file that is used to store information about the updates that are made to the application data folders at the primary site. This information is further used to replicate those updates at the secondary site. This file is also referred to as the replication log.

on-premises system or on-premises application host

A physical computer or virtual machine that exists on the premises of an organization, rather than in the cloud. This system hosts the application that is configured for monitoring and then further configured for recovery.

primary

The system or location that is the source for data replication, where data synchronization is required for recovery in the event of an operational failure. For example, while the application processes requests from the on-premises application host, that system is the primary and the Console host is the secondary.

recovery or disaster recovery (DR)

A feature of Disaster Recovery Orchestrator that enables you to recover application processing in the cloud when your organization's on-premises site becomes unavailable. You can restore application processing back to the on-premises site when it is available again.

You can also perform a planned migration of the application from the on-premises site to the cloud.

recovery administrator (role)

A user who has the privileges to configure an application for recovery, perform recovery operations, and remove the application recovery configuration. A security administrator adds this role for a user to the Privilege Settings tab of the Console UI.

replication log volume

A dedicated volume on the storage that is attached to the systems that participate in the DR solution. This volume is used to store the journal file.

secondary

The system or location that is the destination for data replication, where data synchronization is required for recovery in the event of an operational failure. For example, while the application processes requests from the cloud application host, that Amazon EC2 instance is the primary and the on-premises application host is the secondary.

security administrator

A user who has the privileges to configure the recovery settings and other users for Disaster Recovery Orchestrator Console. The security administrator cannot directly work with application recovery configurations.

takeover

The operation in which application processing is taken over by the cloud application host, when your on-premises application or its host or the site becomes unavailable. A recovery administrator manually triggers Takeover, but the tasks that are involved in the operation are performed automatically in a predefined sequence.

virtualization host

A virtualization management software that creates and manages virtual machines, for example, Microsoft Hyper-V Server or VMware Server.

virtual machine

A software-based computer that is provisioned to run certain processes or provide some specific services, like hosting an application. For example, your on-premises application host can be a VMware or a Hyper-V virtual machine, and your cloud application host is an Amazon EC2 instance.

virtual private network (VPN)

A network connection between your on-premises network and your VPC, which comprises a virtual private gateway and a customer gateway. A virtual private gateway is an AWS component, and the customer gateway is a hardware device or a software application that manages the connections at on-premises site.

Chapter 2

Requirements

This chapter includes the following topics:

System requirements

This section describes the hardware requirements for Disaster Recovery Orchestrator.

Processors

Disaster Recovery Orchestrator requires only a single CPU, and the minimum recommended processing speed is 1 GHz.

The following (or faster) processors are recommended:

- AMD Opteron
- AMD Athlon 64
- Intel Xeon with Intel EM64T support
- Intel Pentium IV with EM64T support

Memory

A minimum of 1 GB of RAM is required to install and use Disaster Recovery Orchestrator.

Storage

Consider the following constraints when you plan or configure the storage to be used for the disaster recovery (DR) solution:

- All the volumes on the on-premises systems as well as the EC2 instances must be NTFS-formatted.

  Note: The recovery of applications depends on the file replication mechanism of Disaster Recovery Orchestrator, which does not work with the FAT or the ReFS file systems.

- GPT disks, dynamic disks, and spanned volumes are not supported on the EC2 instances.
- Only EBS volumes must be used with the EC2 instances.
- Plan the root device volume sizes of the cloud application hosts such that they can accommodate the following components:
  - Operating system
  - Application to be configured for DR
  - Disaster Recovery Orchestrator Client

The following table lists the minimum disk space that is required to install each of the Disaster Recovery Orchestrator components.

Table 2-1 Disaster Recovery Orchestrator disk space requirements

Components	Minimum disk space required
Disaster Recovery Orchestrator Console	700 MB
Disaster Recovery Orchestrator Client	600 MB

The following EBS volumes are required for storage in AWS:

- A root device volume (the operating system disk) for the Console host

  Note: No additional volume must be attached to the Console host.

- For the cloud application host:
  - A root device volume
  - As many application data volumes as required for the configured application

- A separate volume to store the journal file for the application recovery configuration

Software requirements

This topic lists the software that Disaster Recovery Orchestrator requires for successful installation, configuration, and operation.

For the most recent information about the supported software, refer to the software compatibility list (SCL) at:

http://www.symantec.com/docs/TECH225292

Operating systems

You can install and use the Disaster Recovery Orchestrator components on the Windows Server 2008 R2, 2012, and 2012 R2 operating systems. For details about the editions and service packs of these versions that are supported, refer to the SCL.

Note: At the on-premises site, the application may be hosted on any of the supported Windows versions and editions.

In AWS, the available Amazon Machine Images (AMIs) determine the Windows versions and editions that you can use to launch the Amazon EC2 instances. You may create a virtual machine by uploading an image; there is no restriction on the operating system Edition in this case.

However, an on-premises application host and the corresponding cloud application host must run the same operating system version. You may use different editions of the same operating system version, though.

Applications

The following table lists the applications for which Disaster Recovery Orchestrator provides application monitoring and disaster recovery services.

Table 2-2 Disaster Recovery Orchestrator supported applications

Application	Architecture
Microsoft SQL Server 2008 R2	64-bit
Microsoft SQL Server 2012	64-bit
Custom applications and generic services	32-bit, 64-bit

For information about the service packs that are supported with these versions of SQL Server that are supported, refer to the SCL.

Browsers and other software

The following browsers are supported:

- Internet Explorer 9 or later
- Mozilla Firefox 19 or later

The following software is required:

- .NET framework 4.5 is required to install Disaster Recovery Orchestrator components.
- Adobe Flash Player 12 or later is required to use the Disaster Recovery Orchestrator Console UI.
- Your Amazon EC2 instances must use the AWS or the Citrix paravirtualized device drivers. These drivers allow a maximum of 21 volumes to be attached to a Windows instance apart from the root device volume and the removable media.
  If an EC2 instance uses the Red Hat driver, upgrade to the AWS or the Citrix driver. For information about upgrading the paravirtual drivers, see:
  http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/Upgrading_PV_drivers.html

Network and security requirements

The network and security requirements for Disaster Recovery Orchestrator are as follows:

- Virtual private cloud (VPC) and virtual private network (VPN) for disaster recovery (DR)
  - Create an Amazon VPC instance to be used as the failover network where application processing can continue if a failure occurs at the on-premises site. To test the DR configuration for your application without affecting the production environment, use a different subnet within the same VPC for the fire drill network. Alternatively, create a separate VPC in the cloud for performing fire drills.
  - Configure a virtual private gateway in the cloud, a customer gateway at your on-premises site, and the tunnels to establish connectivity between the on-premises and the cloud networks. A VPN connection allows the resources in both the networks to communicate directly and securely.

  See “About configuring VPC and VPN” on page 31.

- Firewall and ports
  If you have firewall, add exceptions for the ports that the applications need to use.
  See "Ports required for Disaster Recovery Orchestrator" on page 23.
- Network protocol usage restriction
  Disaster Recovery Orchestrator does not support the use of IPv6; disable IPv6 on all the systems on which you plan to install the product.
- Domain configuration
  - The DNS server records must be up-to-date at all times so that the Disaster Recovery Orchestrator services do not encounter communication issues.
    See "Domain configuration recommendations" on page 29.
  - Set up reverse DNS and PTR records so that the DNS can map an IP address to a domain name.
- Users and privileges
  Make sure that the existing users have the necessary privileges to install the Disaster Recovery Orchestrator components. Create users with the appropriate privileges so that they can configure applications for monitoring and DR, and perform operations on those configurations.
  See "Users and privileges required for Disaster Recovery Orchestrator" on page 24.
- The following exceptions must be added to the sites that can be accessed from your browser:
  - For the Disaster Recovery Orchestrator Console UI:
    `https://ConsoleHost:14155/draas/login.html`
    Replace the *ConsoleHost* variable with the fully-qualified domain name (FQDN) of the Console host.
  - For managing application monitoring configurations using Disaster Recovery Orchestrator Client:
    `https://ClientHost:5634/vcs/admin/application_health.html`
    Replace the *ClientHost* variable with the FQDNs of the Client hosts.
- The following exceptions must be added to the pop-up blocker on your browser:
  - To enable pop-ups from the Console UI on a system within the private network:

    `ConsoleHost:14155`

    Replace the *ConsoleHost* variable with the FQDN of the Console host.
- To enable pop-ups from the Console UI on a system over the public Internet:

  `ConsoleHostPublicDNSName:14155`

Replace the *ConsoleHostPublicDNSName* variable with the public DNS name of the Console host.

Note: This is applicable only if the Console host is assigned to a security group that allows public access to the resource.

Ports required for Disaster Recovery Orchestrator

Disaster Recovery Orchestrator and its related services need to use some dedicated ports on the systems that participate in the DR solution. If you have configured a firewall, ensure that the firewall settings allow access to the required services and ports.

The following table provides information about the required ports.

Table 2-3 Ports required by Disaster Recovery Orchestrator services

Port	Service	Protocol & Binding	Action Required
5634	Symantec Storage Foundation Messaging Service (`xprtld.exe`)	HTTPS: Bidirectional	Add exception to firewall
14141	Symantec High Availability Engine (`had.exe`)	TCP: Inbound	Ensure availability
14151	Symantec DRaaS Service listens on this port for a shutdown request	TCP: Inbound	Ensure availability
14153	Symantec DRaaS Authentication Service (`vxatd.exe`)	TCP: Inbound	Add exception to firewall
14154	Disaster Recovery Orchestrator Console Database (`dbsrv11.exe`)	HTTPS: Bidirectional	Add exception to firewall
14155	Symantec DRaaS Service (`draasctlsvc.exe`)	TCP: Inbound	Add exception to firewall
14159	Symantec File Replication (`vxrepservice.exe`)	TCP: Bidirectional	Add exception to firewall

Table 2-3 Ports required by Disaster Recovery Orchestrator services *(continued)*

Port	Service	Protocol & Binding	Action Required
49152 – 65535	DCOM Required for ports used by the Symantec File Replication service	TCP: Bidirectional	Add exception to firewall

After adding the necessary exceptions to the firewall, perform the following activities to make the dynamic ports range for DCOM (49152 – 65535) available:

- Make changes to the registry as described in the Microsoft article: http://support.microsoft.com/kb/154596
- Enable the following predefined firewall rules:
  - COM+ Network Access
  - COM+ Remote Administration

Additionally, the following system service ports need to be opened:

25, 53, 67, 88, 123, 135, 137, 138, 139, 389, 443, 445, 464, 636, 1433, 2148, 2535, 3268, 3269, 3389, 5722, 9389

For further information about these ports and their usage, see the Microsoft article:

http://support.microsoft.com/kb/832017

AWS also uses security groups to control inbound and outbound traffic for the instances that you launch in a VPC.

Configure your security groups and assign the Amazon EC2 instances to those groups such that:

- The Disaster Recovery Orchestrator-related services on all the systems that are associated with an application recovery configuration must be able to communicate using the required ports.
- The Apache Tomcat service on the Console host must be able to communicate using the required ports.

Users and privileges required for Disaster Recovery Orchestrator

The following table describes the required users and privileges.

Table 2-4 User and privileges required for Disaster Recovery Orchestrator

Roles	Functions	Privileges
-	Disaster Recovery Orchestrator Console installation	■ Domain user ■ Local administrator on the Console host
-	Disaster Recovery Orchestrator Client installation	■ Local administrator on the Client host ■ May or may not be a domain user
Security administrator	■ Manages the recovery settings on Disaster Recovery Orchestrator Console ■ Manages privilege settings: ■ Designates users as recovery administrators and guests on the on-premises application hosts ■ Edits Admin or Guest privileges for a user ■ Removes user privileges	■ Domain user or domain administrator ■ Local administrator on the Console host
■ Recovery administrator ■ Guest user	Recovery administrators can: ■ Create application monitoring configurations and application recovery configurations. ■ Perform operations on these configurations. ■ Remove these configurations. Guest users can only view the applications, the status of operations that are performed on the application configurations, and the corresponding reports.	Local administrator on the on-premises application host, or the corresponding cloud application host, or both For information about the Console users, refer to the *Symantec Disaster Recovery Orchestrator Administration Guide*.

Table 2-4 User and privileges required for Disaster Recovery Orchestrator *(continued)*

Roles	Functions	Privileges
-	This user's credentials are sought as input when a recovery administrator configures an application for recovery in the cloud (Disaster Recovery Configuration wizard, Virtual Computer Name panel). The Lanman agent uses these credentials to access data and the application on an application host.	■ DNS administrator ■ Local administrator on all the systems that are associated with each application recovery configuration (the on-premises application host, the corresponding cloud application host, and the Console host) For information about the considerations for a DR configuration, refer to the *Symantec Disaster Recovery Orchestrator Administration Guide*.

Requirements for file replication

Disaster Recovery Orchestrator uses file-level replication to ensure that the application data at the on-premises site and at the cloud site is in sync.

For more information about file replication, refer to the *Symantec Disaster Recovery Orchestrator Administration Guide*.

The following requirements must be met for file replication to work:

- If a firewall is enabled on a system that hosts the file replication service, add exceptions to allow traffic across the firewall. These exceptions should include the default file replication port (14159) and any user-configured ports.
- The appropriate security groups must be configured for the VPC in which the EC2 instances for the Console, the application, and the fire drill operation are launched.
- The disks on which application data is stored must be NTFS-formatted. File replication does not work with the FAT or ReFS file systems.
- The following types of files are are not replicated:
  - Reparse points
  - Compressed files
  - Encrypted files

If such files exist in your application data folders, make sure that you manage them appropriately.

- The minimum required journal file size is 1 GB, and the Disaster Recovery Configuration wizard sets this by default. However, Symantec recommends that you set the journal file size to 10 GB so that you do not encounter performance issues.

  You can specify the journal file size at the following locations:

  - On the Replication Settings page of the Disaster Recovery Configuration wizard
    The wizard does not let you proceed with configuring an application for DR if you specify a size less than 1024 MB.
  - On the Settings page of an application on the Console UI
    If you change the size to a smaller value, Disaster Recovery Orchestrator displays an error and does not let you save the change.

- The journal file must *not* be stored at the following locations:

  - The system volume
  - A temporary storage
  - The volume where the application data is stored

  Also, each time a recovery operation is performed on an application, the volume that contains the journal file in the cloud is moved as follows:

  - During the takeover operation, the volume is detached from the Console host and attached to the cloud application host.
  - During the failback operation, the volume is detached from the cloud application host and attached to the Console host.

  Therefore, the journal file for an application configuration at the cloud site must be stored on an independent volume.

Chapter 3

Preparing the cloud environment

This chapter includes the following topics:

Setting up the cloud environment for disaster recovery readiness

Before you can install the Disaster Recovery Orchestrator components, you need to set up the cloud environment.

This topic lists the tasks that are involved in setting up the Amazon Web Services (AWS) environment.

Note: If you have already configured your AWS environment, you may not need to perform the following tasks.

To set up the cloud infrastructure

1. Sign in to the AWS Management Console:

   http://aws.amazon.com/

If you do not already have an AWS account, procure one. To get started, watch the videos that are available at:

https://aws.amazon.com/getting-started/

2. Define Identity and Access Management (IAM) users and generate access keys for the following purposes:
   - To configure and use the appropriate web services that are required for Disaster Recovery Orchestrator.
   - To install and administer the Disaster Recovery Orchestrator components and the application recovery configurations.

   See "About configuring IAM users and permissions" on page 31.
3. Create a virtual private cloud (VPC) and establish a virtual private network (VPN) connection between your on-premises network and the VPC for the following purposes:
   - To create and administer application recovery configurations.
   - To test whether the on-premises applications can be successfully recovered in the cloud.

   See "About configuring VPC and VPN" on page 31.
4. Launch Amazon EC2 Windows instances in the VPC for the following purposes:
   - To host the Disaster Recovery Orchestrator Console components; the Console is used to configure and administer the application recovery configurations.
   - To host the applications and the Disaster Recovery Orchestrator Client components; the Client is used to monitor the applications.

   See "About configuring Amazon EC2 instances" on page 32.

Domain configuration recommendations

Configure a domain controller in the cloud to locally autenticate the users, applications, and services in the cloud. While creating the domain, make the following ports available for the cloud domain controller:

Table 3-1 Ports to be made available for the cloud domain controller

Name	Protocol	Public Port	Private Port	Load-Balanced Set Name
DNS	TCP	53	53	-

Table 3-1 Ports to be made available for the cloud domain controller *(continued)*

Name	Protocol	Public Port	Private Port	Load-Balanced Set Name
LDAP	TCP	389	389	-

These ports are required for Microsoft's site management feature of the domain controller to work.

For more information about the required ports, refer to the Microsoft article:

http://technet.microsoft.com/en-us/library/dd772723(v=ws.10).aspx

---

Note: Symantec recommends that you replicate the on-premises domain controller in the cloud. However, you may choose to create different subnet-based sites and associate them with different domain controllers as per your requirement.

For more information about replicating the domain controller, see:

http://azure.microsoft.com/en-us/documentation/articles/virtual-networks-install-replica-active-directory-domain-controller/

This link is provided only for reference. Even though it points to the Microsoft Azure documentation, the steps for replicating the domain controller in AWS are quite similar.

For information about the other ports required for Disaster Recovery Orchestrator:

See "Ports required for Disaster Recovery Orchestrator" on page 23.

Test the domain controller configuration using the following command:

```
nltest /DSGETDC:DomainName
```

Make sure that the following output string points to the local site:

```
Our Site Name: LocalSiteName
```

If it points to a non-local site, the issue might be one of the following:

- The site-level domain controller configuration is not done.
- The sites are unable to communicate with each other using the aforementioned ports.

To fix this issue, take the following actions:

1. Create a subnet-based site; refer to the Microsoft article:

   http://technet.microsoft.com/en-us/library/cc770372.aspx

2. Enable SYSVOL sharing; refer to the Microsoft article:

   http://support.microsoft.com/kb/947022

3. Reboot the domain controller system.

About configuring IAM users and permissions

To perform any operations in the Amazon Web Services (AWS) cloud, you need to create users and assign them the appropriate permissions. The Amazon Identity and Access Management (IAM) service lets you manage these activities.

To create an IAM user in your AWS account, see:

http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html

For more information about working with IAM users and groups, see:

http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

Users need a valid set of security credentials, called access keys, to perform operations in AWS.

For more information about access keys, see:

http://aws.amazon.com/developers/access-keys/

The Disaster Recovery Orchestrator Console and Client components are installed on Amazon Elastic Cloud Compute (EC2) instances in AWS. Applications that run on Amazon EC2 instances need to be granted access to AWS resources. For example, Disaster Recovery Orchestrator Console sends requests to AWS to perform certain operations, like creating or deleting instances, attaching or detaching volumes, and so on. AWS provides the requested services if the Console has the appropriate permissions.

To use IAM roles to grant applications access to the cloud resources, see:

http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

About configuring VPC and VPN

Disaster Recovery Orchestrator requires you to use the Amazon virtual private cloud (VPC) service within which to manage your cloud resources.

To implement the disaster recovery (DR) solution, you need to configure the following network-related entities:

- Create an Amazon VPC, its subnets, and the security groups as per your network security requirements.

To get started with this exercise, see:
http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/ExerciseOverview.html
For more information about Amazon VPC, see:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

- Establish a VPN connection between the VPC and your on-premises network. A VPN connection consists of a virtual private gateway and a customer gateway.
  For information about the manual process, see:
  http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

  Alternatively, you may use the VPC creation wizard to set up the virtual private gateway for the scenario that is relevant to you:

  - VPC with public and private subnets and hardware VPN access
    http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario3.html
  - VPC with a private subnet only and hardware VPN access
    http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario4.html

  For information about customer gateways, see:
  http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html

Keep the following information ready before you perform network configurations in the cloud:

- The IP addresses and subnet ranges of the on-premises network
- The IP addresses and subnet ranges that you plan to use in the cloud
- An external IP address for the virtual private network (VPN) device

About configuring Amazon EC2 instances

Disaster Recovery Orchestrator requires you to use the Amazon EC2 service for managing the cloud computing resources. Disaster Recovery Orchestrator only supports EC2 Windows instances, which are virtual machines running Windows Server that you can create and manage in the cloud.

To implement the DR solution, you need to launch the following Windows instances:

- A dedicate Windows instance to manage the DR configurations. You install and configure Disaster Recovery Orchestrator Console on this instance, and then refer to it as the Console host. It manages the DR operations and the replication

that is required to keep the data in sync between the on-premises and the cloud sites.

- A Windows instance corresponding to each on-premises system that hosts an application that you want to configure for DR. You install and configure Disaster Recovery Orchestrator Client and the appropriate application on each such instance, and then refer to them as the cloud application hosts.

Note: A cloud application host must run the same operating system version as the corresponding on-premises application host. You may use different editions of the same operating system version, though.

Consider the system requirements and the purpose of the Windows instance (Console host or cloud application host) to identify the type of instance to be used.

For information about the system requirements:

See "System requirements" on page 18.

For information about the types of EC2 instances, see:

http://aws.amazon.com/ec2/instance-types/

To launch a Windows instance, perform the following tasks:

1. If you haven't already done so, prepare to use the Amazon EC2 service.

   To get the basic setup ready, see:

   http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/get-set-up-for-amazon-ec2-windows.html

2. Launch a Windows instance and connect to it.

   To get started with this exercise, see:

   http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/EC2Win_GetStarted.html

   Optionally, you can create an Amazon EBS-backed Windows AMI based on a Windows instance that you launch and customize.

   To create an EBS-backed Windows AMI, see:

   http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/Creating_EBSbacked_WinAMI.html

Note: If you use the same AMI to launch multiple instances, change the Computer Name of the Windows instance so that it is unique in the VPC.

For more information about EC2, see:

http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/IntroWindowsUserGuide.html

For more information about EBS, see:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html

Chapter 4

Installing and configuring the product

This chapter includes the following topics:

About installing Disaster Recovery Orchestrator components

Installing Disaster Recovery Orchestrator involves the following tasks:

- Install Disaster Recovery Orchestrator Console on an Amazon EC2 instance. This instance would then function as a controller for the disaster recovery (DR) activities.
- Optionally, configure a backup of the Console components. This adds resiliency to Disaster Recovery Orchestrator Console in case it fails due to corruption in the operating system or the applications.
- Install Disaster Recovery Orchestrator Client on the on-premises application host and on the corresponding cloud application host. The application will be available on the cloud application host when the on-premises application or its host becomes unavailable.

Considerations for installing Disaster Recovery Orchestrator Console

Consider the following before installing Disaster Recovery Orchestrator Console:

- Internet connectivity is required to access the cloud environment.
- An Amazon Web Services (AWS) account is required to sign in to the AWS Management Portal.
- IAM users with the required access keys and permissions for the cloud authentication.
  See "About configuring IAM users and permissions" on page 31.
- Disaster Recovery Orchestrator Console can be installed only on an Amazon EC2 instance with one of the supported platforms.
  See "Software requirements" on page 20.
- The Amazon EC2 instance on which you install Disaster Recovery Orchestrator Console is referred to as the Console host. The installer must be launched locally on this instance. Remote installation of the Console components is not allowed.
- An external disk should not be manually attached to the Console host.

  Note: Only Disaster Recovery Orchestrator should manage the addition or removal of all storage devices on the Console host.

- The following user configurations are expected:
  - The user who installs Disaster Recovery Orchestrator Console must be a valid domain user and a member of the local Administrators group.
  - The User Access Control (UAC) feature of Windows must be disabled.
- A separate subnet or a separate VPC is required to configure fire drills. It must not be a production network.
  Optionally, you can create this fire drill network later, and add its details to the Settings page of Disaster Recovery Orchestrator Console.
- Amazon EC2 instances running Windows Server 2008 R2 must have .NET 4.5 installed. It is available by default with Windows Server 2012 and 2012 R2.

Installing Disaster Recovery Orchestrator Console

Install Disaster Recovery Orchestrator Console on an Amazon EC2 instance that you plan to designate as a controller for the disaster recovery (DR) activities.

To install Disaster Recovery Orchestrator Console using the installation wizard

1. In Windows Explorer, browse to the Disaster Recovery Orchestrator software package directory, and double-click the `Setup.exe` file.
2. On the Symantec Product Installer screen, click **Install Disaster Recovery Orchestrator Console** to launch the installation wizard.
3. On the Welcome panel, review the list of prerequisites, make sure that they are met, and then click **Next**.
4. On the License panel, read the license terms, select **I accept the terms of License Agreement**, and click **Next**.

   The **Participate in the Symantec Product Improvement Program by submitting system and usage information anonymously** check box is selected by default. The Product Improvement Program allows the product installer to collect installation, deployment, and usage data, and anonymously submit it to Symantec. This information helps identify how customers deploy and use the product. If you do not want to participate in the product improvement program, clear this check box.
5. On the System Validation panel, verify the following:
   - The current system appears in the System field.
     Remote installation is not allowed.
   - The default installation directory is: `C:\Program Files\Veritas`. You can customize this location either by entering the path manually or by clicking **Browse...** and navigating to the desired location.
   - The wizard performs certain validation checks on the system and notes the details in the Verification Status field.
     If the system fails the validation checks, the wizard does not proceed with the installation. Review the details, rectify the issue, and then click **Re-Verify** to re-initiate the validation checks for this system.

   Click **Next**.

6 On the Cloud Authentication panel, provide the appropriate user access keys and click **Next**.

The wizard uses this information to authenticate the user credentials in the cloud environment.

Note: The cloud region indicates your cloud network location, and is selected automatically.

For information about configuring IAM users and their access keys:

See "About configuring IAM users and permissions" on page 31.

7 On the Fire Drill Network panel, provide the following information:

- Select the Amazon virtual private cloud (VPC) to be used for performing fire drills. If you select the same VPC in which the current Amazon EC2 instance is launched, make sure to select a different subnet.
- Select the subnet to be used.
- Select one or more security groups for this Amazon EC2 instance.

Alternatively, you may specify these details later on the DR Settings page of Disaster Recovery Orchestrator Console.

Click **Next**.

For information about configuring networks for DR in AWS:

See "About configuring VPC and VPN" on page 31.

8 On the Domain Authentication panel, provide the credentials of the Active Directory user for authentication on the Windows domain.

9 On the Pre-install Summary panel, review the summary.

If you want to save this information for future reference, click **Save Report**. The default file name is `PreInstallReport`, and you can save it in a few different formats.

If any issues occur, review the log for details, and rectify the issues before proceeding.

Click **Next**.

10 On the Installation panel, review the progress of the installation.

If the installation fails, the details pane displays the status accordingly. You may click **Cancel** to exit the wizard. Perform the necessary troubleshooting tasks, and relaunch the wizard to complete the installation.

After the installation succeeds, click **Next** to begin the post-installation tasks.

11 On the Post-install Summary panel, review the progress of the post-installation tasks and the summary of the all the changes made by the installer.

As part of the post-installation tasks, the wizard applies the `SDRO_Role=SDRO_Console_Instance` tag to the Console host. For more information about tagging Amazon EC2 resources, see:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

If you want to save this information for future reference, click **Save Report**. The default file name is `PostInstallReport`, and you can save it in a few different formats.

If the installation fails, refer to the log file for details. You may have to reinstall the software.

Click **Finish** to exit the wizard.

Adding resiliency to Disaster Recovery Orchestrator Console

Disaster Recovery Orchestrator Console may be susceptible to corruption in the operating system or the applications. Symantec recommends that you back up the Console configuration to make it less susceptible to corruption. The following utilities enable you to add resiliency to the Console:

- `backup_console.pl`

  This utility performs the following actions:

  - Takes a one-time backup of the Disaster Recovery Orchestrator database and the authentication configuration
  - Schedules the periodic backups of the Disaster Recovery Orchestrator database

  You can specify the frequency and the location of the backup.

- `restore_console.pl`

  This utility restores the authentication configuration and database files, which includes attaching the appropriate data disks to the new Console host.

Note: You may perform the tasks that are listed in this topic after you have configured the applications for monitoring and recovery in the cloud. However, if you want to make the Console less susceptible to corruption right from the beginning, perform these tasks immediately after the installation.

To add resiliency to Disaster Recovery Orchestrator Console

1. Make a note of the following items:
   - Host name of the Amazon EC2 instance that is the Console host
   - Names of the storage volumes that are attached to the Console host
     The root device volume can be deleted if the Console host needs to be re-created. The volumes on which the journal file and the application data are stored must be retained.
2. Schedule a backup.

   Scheduling is a one-time activity that you can perform immediately after installing Disaster Recovery Orchestrator Console or after configuring applications for recovery.

   Refer to the following procedure.

   To schedule a backup
3. If a Disaster Recovery Orchestrator Console failure occurs, perform the following tasks sequentially:
   - If the replication volume and the application data volumes are attached to the Console host, ensure that their **Delete On Termination** property is set to `False`, which is the default.
     If this property is set to `True` for a volume, the volume is deleted when the instance to which it is attached is terminated.
   - Terminate the Console host from the AWS Management Portal.
     You may also delete the root device volume that was attached to the Console host.
     For information about terminating the EC2 instance, see:
     http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/terminating-instances.html
   - Launch an Amazon EC2 instance and change its host name to the value that was used for the previous Console host.

- Install Disaster Recovery Orchestrator Console on the newly launched instance.

4 Restore the Disaster Recovery Orchestrator Console configuration using the backup.

Refer to the following procedure.

To restore the backup

To schedule a backup

1 Log on to the Console host.

2 Open the command prompt at the Console installation folder. For example:

```
C:\Program Files\Veritas\draasconsole\bin
```

3 Run the following command:

```
"C:\Program Files\Veritas\vrtssfmh\bin\perl.exe" backup_console.pl
/PATH NetworkLocation /USERNAME UserName [/PASSWORD Password]
[/PERIOD Period]
```

Here, the parameters take the following values:

- *NetworkLocation* is the location where the data is backed up.

  Note: This location must be a file share on a Windows instance within the same domain.

- *UserName* and *Password* are the credentials of any user who has access permissions on the network location and the local system.
- *Period* is the duration (in hours) after which database is backed up; the default is 2 hours.

This utility stops the Symantec DRaaS Console Database Service, the Symantec DRaaS Service, and the Symantec DRaaS Authentication Service. It starts these services again immediately after the backup task is complete. This one-time activity is performed only at the first instance of the backup. These services are not stopped during the subsequent database backups.

To stop the periodic backup

1 Open the Task Scheduler window.

2 Expand **Task Scheduler Library > Symantec**.

3 Select the **DRaaS_Console_Backup** task, and click **Delete** from the Actions menu.

The Console is no longer backed up.

4 If you want to schedule the periodic backups again, run `backup_console.pl`.

Refer to the following procedure.

To schedule a backup

To restore the backup

1 Log on to the Console host.

2 Open the command prompt at the Console installation folder. For example:

```
C:\Program Files\Veritas\draasconsole\bin
```

3 Run the following command:

```
"C:\Program Files\Veritas\vrtssfmh\bin\perl.exe"
restore_console.pl NetworkLocation
```

Here, *NetworkLocation* is the location where the data is backed up.

Considerations for installing Disaster Recovery Orchestrator Client

Before you begin to install Disaster Recovery Orchestrator Client, consider the following:

- The user who performs the installation must have local administrator privileges. In case of remote installations, the user must have local administrator privileges on all the selected systems.
- On the systems where you plan to install these components:
  - The User Access Control (UAC) feature of Windows must be disabled.
  - The .NET 4.5 framework must be installed.
  - IPv6 must not be configured.
  - If the system is an Amazon EC2 instance, only EBS volumes must be attached.

Installing Disaster Recovery Orchestrator Client

For each application that you want to configure for disaster recover (DR), install Disaster Recovery Orchestrator Client on the following systems:

- The on-premises system that hosts the application (on-premises application host)
- The Amazon EC2 instance (cloud application host) where the application will be available when the on-premises application or its host becomes unavailable

To install Disaster Recovery Orchestrator Client using the installation wizard

1. In Windows Explorer, browse to the Disaster Recovery Orchestrator software package directory, and double-click the `Setup.exe` file.
2. On the Symantec Product Installer screen, click **Install Disaster Recovery Orchestrator Client** to launch the installation wizard.
3. On the Welcome panel, review the list of prerequisites, make sure that they are met, and then click **Next**.
4. On the License panel, read the license terms, select **I accept the terms of License Agreement**, and click **Next**.

   The **Participate in the Symantec Product Improvement Program by submitting system and usage information anonymously** check box is selected by default. The Product Improvement Program allows the product installer to collect installation, deployment, and usage data, and anonymously submit it to Symantec. This information helps identify how customers deploy and use the product. If you do not want to participate in the product improvement program, clear this check box.
5. On the System Validation panel, select the systems and their installation directories as follows:

   The `localhost` information is populated by default; you can remove it if you do not want to install Disaster Recovery Orchestrator Client on it. You can install the product remotely on other systems in the domain. You can select more systems in the following ways:

   - In the **System Name or IP** field, type the name or IP address of a system, and click **Add**.

     Note: Disaster Recovery Orchestrator does not support Internet Protocol version 6 (IPv6), so do not add a system that uses IPv6.

   - Alternatively, click **Browse...** to select the systems.
     The systems that belong to the domain to which you have logged on are listed in the **Available Systems** list. Select one or more systems and move them to the **Selected Systems** list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks, and notes the details in the Verification Details field. To review the details of a particular system, select it from the list on the left.

The default installation directory is: `C:\Program Files\Veritas`. For each system, you can customize this location either by entering the path manually or by clicking **Browse...** and navigating to the desired location.

To use the customized location for to multiple systems, click **Apply Install Options to Multiple Systems...**. On the dialog box that appears, select the systems on which to customize the location, and then click **OK**.

Unless all the selected systems pass the validation checks, the wizard does not proceed with the installation. For each system that might have failed the validation checks, review the details, rectify the issue, and then click **Re-Verify**.

When all the systems have been verified, click **Next** to proceed with the installation.

6 On the Pre-install Summary panel, review the summary and click **Next**.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PreInstallReport`, and you can save it in a few different formats.

If any issues occur, review the log for details, and rectify the issues before proceeding.

7 On the Installation panel, review the progress of the installation.

If the installation fails, the details pane displays the status accordingly. You may click **Cancel** to exit the wizard. Perform the necessary troubleshooting tasks, and relaunch the wizard to complete the installation.

After the installation succeeds, click **Next** to begin the post-installation tasks.

8 On the Post-install Summary panel, review the progress of the post-installation tasks and the summary of the all the changes made by the wizard.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PostInstallReport`, and you can save it in a few different formats.

If the installation fails, refer to the log file for details. You may have to reinstall the software.

Click **Finish** to exit the wizard.

Chapter 5

Repairing the product installation

This chapter includes the following topics:

Considerations for repairing a Disaster Recovery Orchestrator installation

Repairing an installation restores it to its original state, which includes replacing missing or corrupt files, shortcuts, and registry entries. Use the Disaster Recovery Orchestrator installer to perform a repair operation.

Before you begin to repair an installation, consider the following:

- Before repairing the Disaster Recovery Orchestrator installation, repair the **Veritas Operations Manager (Host Component)** on the system.
- The installer uses the logged-on user account context to perform the repair operation. Verify that the logged-on user has local administrator privileges on the system where you want to repair the installation.
- If you have configured application monitoring, it may be temporarily suspended while the installer performs the reparation tasks. Therefore, the Health View may not display the most current status of the configured application during this time.
- If you have configured an application disaster recovery (DR), it may be temporarily suspended while the installer performs the reparation tasks.

Therefore, the Disaster Recovery Orchestrator Console UI may not display the most current status of the DR configuration during this time.

- Any ongoing file replication activities must be paused before you repair the installation. You must also resume the activities after the repair operation completes successfully.
- You can repair a local installation only. Repairing an installation remotely is not supported.
- You cannot repair a failed installation. The Repair option is available only for an installation that has completed successfully.
- While repairing the product installation, you cannot modify the installation options.

Repairing a Disaster Recovery Orchestrator Console installation

The Disaster Recovery Orchestrator Console installer runs in the Repair mode to restore the installed components to their original state.

To repair a Disaster Recovery Orchestrator Console setup

1. Pause any ongoing file replication activities by running the following commands sequentially on the Console host:
   - `vxfradmin -viewconfig`
     A list of all the replicated file group (RFG) names, each associated with an application recovery configuation, is displayed.
   - `vxfradmin -pauserep` *`RFGName`*
     Run this command for each RFG to pause the file replication.
2. In the Programs and Features window, select Symantec Disaster Recovery Orchestrator 6.1.1 Console, and click **Change** to launch the installer.
3. On the Mode Selection panel, the **Repair** option is selected by default. Click **Next**.
4. On the System Validation panel, verify the following:
   - The current system appears in the System field.
     Remotely repairing an installation is not allowed.
   - The location in the Install Directory field is correct.
     You cannot specify a different location when repairing an installation.
   - The wizard performs certain validation checks on the system and notes the details in the Verification Status field.

If the system fails the validation checks, the wizard does not proceed with the repair operation. Review the details, rectify the issue, and then click **Re-Verify** to re-initiate the validation checks for this system.

Click **Next**.

5 On the Cloud Authentication panel, provide the appropriate user access keys and click **Next**.

Click **Next**.

The wizard uses this information to authenticate the user credentials in the cloud environment.

Note: The cloud region indicates your cloud network location, and is selected automatically.

For information about configuring IAM users and their access keys:

See "About configuring IAM users and permissions" on page 31.

6 On the Firedrill Network panel, provide the following information:

- Select the Amazon virtual private cloud (VPC) to be used for performing fire drills. If you select the same VPC in which the current Amazon EC2 instance is launched, make sure to select a different subnet.
- Select the subnet to be used.
- Select one or more security groups for this Amazon EC2 instance.

Alternatively, you may specify these details later on the DR Settings page of Disaster Recovery Orchestrator Console.

Click **Next**.

For information about configuring networks for DR in AWS:

See "About configuring VPC and VPN" on page 31.

7 On the Domain Authentication panel, provide the credentials of the Active Directory user for authentication on the Windows domain.

Note that the user must not be a member of the local Administrators group.

8 On the Pre-install Summary panel, review the summary.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PreRepairReport`, and you can save it in a few different formats.

If any issues occur, review the log for details, and rectify the issues before proceeding.

Click **Next**.

9 On the Installation panel, review the progress of the repair operation.

If the repair operation fails, the wizard displays the status accordingly; click **Cancel** to exit the wizard. Then, perform the necessary troubleshooting tasks, and relaunch the wizard to complete the operation.

After the repair operation succeeds, click **Next** to perform the post-repair tasks.

10 On the Post-install Summary panel, review the progress of the post-repair tasks and the summary of the all the changes made by the wizard.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PostRepairReport`, and you can save it in a few different formats.

If the repair operation fails, refer to the log file for details. You may have to reinstall the software.

Click **Finish** to exit the wizard.

11 Resume any file replication activities that were paused by running the following command on the Console host:

```
vxfradmin -resumerep RFGName
```

Run this command to resume the file replication for each RFG for which you paused the activity earlier.

Repairing a Disaster Recovery Orchestrator Client installation

The Disaster Recovery Orchestrator Client installer runs in the Repair mode to restore the installed components to their original state.

To repair a Disaster Recovery Orchestrator Client setup

1 Pause any ongoing file replication activity by running the following commands sequentially on the local system:

- `vxfradmin -viewconfig`

If an application on the system is configured for disaster recovery (DR), the corresponding replicated file group (RFG) name is displayed.

- `vxfradmin -pauserep` *`RFGName`*
  Run this command to pause the file replication.

2 In the Programs and Features window, select Symantec Disaster Recovery Orchestrator 6.1.1 Client, and click **Change** to launch the installer.

3 On the Mode Selection panel, the **Repair** option is selected by default. Click **Next**.

4 On the System Validation panel, verify the following:

- Only the current system appears in the systems list.
  Remotely repairing an installation is not allowed.
- The installed options for this system are uneditable.
  You cannot specify a different location or a license key when repairing an installation.
- The wizard performs certain validation checks on the system and notes the details in the Verification Details field.
  If the system fails the validation checks, the wizard does not proceed with the repair operation. Review the details, rectify the issue, and then click **Re-Verify** to re-initiate the validation checks for this system.

Click **Next**.

5 On the Pre-install Summary panel, review the summary.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PreRepairReport`, and you can save it in a few different formats.

If any issues occur, review the log for details, and rectify the issues before proceeding.

Click **Next**.

6 On the Installation panel, review the progress of the repair operation.

If the repair operation fails, the wizard displays the status accordingly; click **Cancel** to exit the wizard. Then, perform the necessary troubleshooting tasks, and relaunch the wizard to complete the operation.

After the repair operation succeeds, click **Next** to perform the post-repair tasks.

7 On the Post-install Summary panel, review the progress of the post-repair tasks and the summary of the all the changes made by the wizard.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PostRepairReport`, and you can save it in a few different formats.

If the repair operation fails, refer to the log file for details. You may have to reinstall the software.

Click **Finish** to exit the wizard.

8 Resume any file replication activity that was paused by running the following command on the local system:

```
vxfradmin -resumerep RFGName
```

Run this command to resume the file replication that you paused earlier.

Chapter 6

Uninstalling the product

This chapter includes the following topics:

About uninstalling Disaster Recovery Orchestrator components

Before you uninstall the Disaster Recovery Orchestrator components, consider the following:

- If application monitoring is configured on the system, you must remove the configuration.
  If the application is also configured for DR, you must remove the DR configuration.
  Use the Unconfigure menu of the Console UI to remove an application recovery configuration. Doing so also removes the corresponding monitoring configuration from the Client hosts.
  For more information, see the *Symantec Disaster Recovery Orchestrator Administration Guide*.
- The installer uses the logged-on user account context for uninstallation. Verify that the logged-on user has local administrator privileges on the system where you want to uninstall the product.
- Remote uninstallation is not supported.
- The Application Information service must be running on all the systems participate in the DR solution. To start the service type the following at the command prompt:

  ```
  net start appinfo
  ```

Note: Do not uninstall Veritas Operations Manager (Host Component) before uninstalling Disaster Recovery Orchestrator. This component is shared among multiple Symantec products that may be installed on the same system. Uninstall Veritas Operations Manager (Host Component) only after all Symantec products, including Disaster Recovery Orchestrator, are uninstalled from the system.

Uninstalling Disaster Recovery Orchestrator Client

To remove Disaster Recovery Orchestrator Client completely from your setup, perform this operation on each system where it was installed. Unlike the installation, you cannot perform an uninstallation on remote systems.

To uninstall Disaster Recovery Orchestrator Client

1. In the Programs and Features window, select Symantec Disaster Recovery Orchestrator 6.1.1 Client, and click **Uninstall** to launch the installer.
2. On the Welcome panel, review the prerequisites, make sure that they are met, and then click **Next**.
3. On the System Validation panel, perform the following actions:
   - Verify that the current system appears in the System field. Remote uninstallation is not allowed.
   - Verify that the correct path appears in the Install Directory field.
   - The wizard performs certain validation checks on the system and notes the details in the Verification Details field. If the system fails the validation checks, the wizard does not proceed with the uninstallation. Click **OK** to close the message box that appears.
   - Review the details, rectify the issue, and then click **Re-Verify**.
   - When the system has been verified, click **Next**.
4. On the Pre-install Summary panel, review the summary.

   If you want to save this information for future reference, click **Save Report**. The default file name used is `PreUninstallReport`, and you can save it in a few different formats.

   If any issues occur, review the log for details, and rectify the issues before proceeding.

   Click **Next**.

5 On the Uninstallation panel, review the progress of uninstallation.

If the uninstallation fails, the wizard displays the status accordingly; click **Cancel** to exit the wizard. Then, perform the necessary troubleshooting tasks, and relaunch the wizard to complete the operation.

After the uninstallation succeeds, click **Next** to perform the post-uninstallation tasks.

6 On the Post-install Summary panel, review the progress of the post-uninstallation tasks and the summary of the all the changes made by the wizard.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PostUninstallReport`, and you can save it in a few different formats.

If the uninstallation fails, refer to the log file for details. You may have to attempt this operation again, or manually clean up uninstallation.

Click **Finish** to exit the wizard.

Uninstalling Disaster Recovery Orchestrator Console

Uninstall Disaster Recovery Orchestrator Console from the Console host in the cloud.

To uninstall Disaster Recovery Orchestrator Console

1 In the Programs and Features window, select Symantec Disaster Recovery Orchestrator 6.1.1 Console, and click **Uninstall** to launch the installer.

2 On the Welcome panel, review the prerequisites, make sure that they are met, and then click **Next**.

3 On the System Validation panel, perform the following actions:

- Verify that the current system appears in the System field. Remote uninstallation is not allowed.
- Verify that the correct path appears in the Install Directory field.
- The wizard performs certain validation checks on the system and notes the details in the Verification Status field. If the system fails the validation checks, the wizard does not proceed with the uninstallation. Click **OK** to close the message box that appears.
- Review the details, rectify the issue, and then click **Re-Verify**.
- When the system is validated, click **Next**.

4 On the Pre-uninstall Summary panel, review the summary.

If you want to save this information for future reference, click **Save Report**. The default file name used is `PreUninstallReport`, and you can save it in a few different formats.

If any issues occur, review the log for details, and rectify the issues before proceeding.

Click **Next**.

5 On the Uninstallation panel, review the progress of uninstallation.

If the uninstallation fails, the wizard displays the status accordingly; click **Cancel** to exit the wizard. Then, perform the necessary troubleshooting tasks, and relaunch the wizard to complete the operation.

After the uninstallation succeeds, click **Next** to perform the post-uninstallation tasks.

6 On the Post-uninstall Summary panel, review the progress of the post-uninstallation tasks and the summary of the all the changes made by the wizard.

As part of the post-uninstallation tasks, the wizard removes the `SDRO_Role=SDRO_Console_Instance` tag from the Console host. For more information about tagging Amazon EC2 resources, see:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

If you want to save this information for future reference, click **Save Report**. The default file name used is `PostUninstallReport`, and you can save it in a few different formats.

If the uninstallation fails, refer to the log file for details. You may have to attempt this operation again, or manually clean up uninstallation.

Click **Finish** to exit the wizard.

Appendix A

Troubleshooting

This appendix includes the following topics:

Disaster Recovery Orchestrator logging

Disaster Recovery Orchestrator provides the following logging information.

Installation logs

Disaster Recovery Orchestrator installer logs contain details about the installation tasks and the overall progress status. These logs are useful for identifying installation-related issues.

The installer creates the log directory as soon as you launch the wizard. The log files are located at `%AllUsersProfile%\Veritas\VPI\log\`. The `%AllUsersProfile%` environment variable expands to `C:\ProgramData`.

Console logs

The Disaster Recovery Orchestrator Console logs are located at `%AllUsersProfile%\symantec\draasconsole\Logs`.

The Console logs are written to the `SDROConsole.log` file.

The components of the Console logs are as follows:

- Timestamp
  The date and time the message was generated
- Duration

The number of milliseconds elapsed between the construction of the layout and the creation of the logging event

- Thread
  The name of the thread that generated this logging event
- Priority
  Levels in the increasing order of priority: TRACE, DEBUG, INFO, WARN, ERROR, and FATAL
- Object
  The logger object, which the fully-qualified class name of the caller that issues the logging request
- Message
  The actual message that was generated by the thread

Additionally, the Disaster Recovery Orchestrator UI components create logs that are available only for the duration of their existence. For example, the Console UI logs are available only as long as you are signed in and the session is active. These logs are lost after the session ends. Similarly, the Disaster Recovery Configuration wizard creates a log file that is available from within the wizard. This information is lost when you exit the wizard.

Agent logs

The agent logs are located at `%vcs_home%\log`. The `%vcs_home%` environment variable points to the product home directory, typically, `C:\Program Files\Veritas\Cluster Server`.

The components of the agent logs are as follows:

- Timestamp
  The date and time the message was generated
- Mnemonic
  The string ID that represents the product, for example, SDRO
- Severity
  Levels in the increasing order of severity: INFO, NOTICE, WARNING, ERROR, and CRITICAL
- UMI
  A unique message ID
- Message
  The actual message generated by the agent

Collecting Disaster Recovery Orchestrator logs

Disaster Recovery Orchestrator provides the `hagetcf` utility, which you can use to collect logs. This utility retrieves detailed diagnostic information about your application monitoring and recovery configurations. You can use this information to troubleshoot configuration-related issues. You can also share these logs with Symantec Technical Support for further troubleshooting.

The `hagetcf` utility is available at the following locations:

- On the Console host (the Amazon EC2 instance where Disaster Recovery Orchestrator Console is installed):

  InstallDir`\draasconsole\bin`

  Here, *InstallDir* is the Disaster Recovery Orchestrator Console installation directory, typically, `C:\Program Files\Veritas`.

- On a system where Disaster Recovery Orchestrator Client is installed:

  `%vcs_home%\bin`

  The `%vcs_home%` environment variable points to the product home directory, typically, `C:\Program Files\Veritas\Cluster Server`.

To collect Disaster Recovery Orchestrator logs

1. On the Console host, navigate to the location where the utility is installed.

   This step is not required on Disaster Recovery Orchestrator Client systems.

2. Collect the logs using the following command:

   `hagetcf [-`*Option*`]`

   You can limit the diagnostic information to specific components using the various available options.

   Use the `-?` or `-help` option to view the command's usage information.

   Note: If you do not specify any options, the command retrieves diagnostic information with the options: `-app`, `-sys`, `-ha`, `-log`, `-lock`, `-conf`, `-state`,`-islog`, and `-trigger`. On a Console host, it also includes the logs for the Disaster Recovery Orchestrator Console component.

   By default, `hagetcf` writes the output to the following locations:

   - On the Console host:

     `%AllUsersProfile%\Symantec\hagetcf\`*mmdd_hhmm*

     The `%AllUsersProfile%` environment variable points to the common program data location, typically, `C:\ProgramData`.

- On Disaster Recovery Orchestrator Client systems:
  %vcs_home%\hagetcf*mmdd_hhmm*

The *mmdd_hhmm* folder name indicates the date and time when the logs were collected, for example: `C:\Program Files\Veritas\Cluster Server\hagetcf\0428_1520`. The folder contains several subfolders and log files, which represent various components.

Disaster Recovery Orchestrator deployment issues and solutions

This section lists the issues that you might encounter when installing, repairing, or uninstalling the Disaster Recovery Orchestrator components. It also describes the tasks that you can perform to work around these issues.

The Console installer fails to discover the cloud subscription

Even though you provide all the correct input values on the Cloud Authentication panel, the Console installer may still fail to validate the credentials.

Some of the possible reasons could be as follows:

- The internet is not accessible from the Amazon EC2 instance. In this case, contact your system administrator.
- The required cloud services are not working properly. In this case, check the status of the cloud services at the following location:
  http://status.aws.amazon.com/
  You need to wait for the services to be restored before proceeding with the installation. Contact AWS Support for further information.

'ERROR: Could not connect to server (err=167)' occurs during post-install configuration of Disaster Recovery Orchestrator Console

If you encounter this issue, take the following actions:

- Make sure that the Symantec Storage Foundation Messaging Service (`xprtld`) service is running on the system.
- If the `xprtld` service is running, check your domain controller configuration using the following command:
  `nltest /DSGETDC:`*`DomainName`*
  Make sure that the following output string points to the local site:
  `Our Site Name:` *`LocalSiteName`*

If it points to a non-local site, fix the domain controller configuration as per Symantec recommendations.

See "Network and security requirements" on page 21.

Uninstallation of Console or Client fails if the relevant services are not stopped

A Disaster Recovery Orchestrator Console or Client uninstallation may fail if any of the following services are not stopped:

- Symantec DRaaS Authentication Service
- Symantec DRaaS Console Database Service
- Symantec DRaaS Service
- Symantec File Replication

The Uninstallation panel displays a message that mentions the services that have not stopped as expected.

Workaround

Perform the following tasks:

1. Open the Services window, select the services that need to be stopped, and select **Action > Stop**.
2. On the Uninstallation panel, click **Back** and click **Next** again to proceed with the uninstallation.

   Alternatively, exit the wizard and launch the uninstallation program again from the Programs and Features window.

Index

T

U

W