Veritas™ Resiliency Platform 2.0: Solutions for Applications

VERITAS™

Veritas Resiliency Platform: Solutions for Applications

Last updated: 2016-09-13

Document version: 2.0 Rev 1

Legal Notice

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

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. 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:

doc.feedback@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

Overview of Resiliency Platform

This chapter includes the following topics:

About Veritas Resiliency Platform

Veritas Resiliency Platform offers a unified solution that helps you proactively maintain business uptime across private, public, and hybrid clouds. Resiliency Platform gives you complete automation for all resiliency operations involving the virtual machines, applications, and multi-tier business-services in your data center. It safeguards the current technology investments by plugging into your existing environments and infrastructure.

For data replication, you can use the Veritas Data Mover or any third-party solution that is supported by Veritas Resiliency Platform. For a list of supported vendors and products, see *Veritas Resiliency Platform Hardware and Software Compatibility Guide*.

Resiliency Platform has the following core capabilities:

Security and Compliance	Veritas Resiliency Platform provides enhanced data encryption (for data-in-flight and data-at-rest) as well as choice of data residency.
Predictability	Customers can predictably meet critical business Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs).
Compliance	Customers can prove compliance to internal and external business continuity mandates with audit reporting and non-disruptive, real-time disaster recovery testing.
Automation	Customers get complete automation for all resiliency operations including recovery run books, and start and stop recovery orchestration for multi-tier applications. This reduces risk of downtime from human error.
Flexibility	Customers get the flexibility to keep their existing infrastructures and can innovate on their terms, with the flexibility that Resiliency Platform provides, to enable workload migration across sites and even to the cloud.

About disaster recovery using Resiliency Platform

A comprehensive disaster recovery strategy ensures that your mission-critical IT functions can continue during and after a disaster and any unforeseen risk can be mitigated to the extent possible.

Veritas Resiliency Platform lets you perform disaster recovery operations on your critical IT services. This section introduces you to the key features of Resiliency Platform:

- Monitoring of data center assets - storage, virtual machines, and applications.

- Ability to group your virtual machines or applications in resiliency groups based on your production environment and business needs.

- Making business services more resilient by providing the ability to perform disaster recovery operations on virtual machines and applications. For example, migrate and take over.

- Ability to replicate data from virtual machines on source data centers to target data centers using Resiliency Platform Data Mover integrated with VMware API I/O filtering framework or array-based replication technologies provided by array vendors.

- Resiliency plan (a sequential execution of predefined steps) to automate site-level recovery operations on your IT infrastructure in the event of downtime.

- Auto-discovery and real-time tracking for recovery objectives.

- Ability to perform non-disruptive testing (rehearsal) on your virtual machines and applications to ensure that your infrastructure is adequately prepared for protection in the event of disaster.

- Reporting capabilities providing details about resiliency health of applications and virtual machines.

See "Understanding the role of resiliency groups in disaster recovery operations" on page 56.

About Resiliency Platform features and components

The following is a brief introduction to Veritas Resiliency Platform key components and their relationships. Administrators responsible for deploying and configuring the product need to understand these in more detail.

resiliency domain	The logical scope of a Resiliency Platform deployment.
	It can extend across multiple data centers.
Resiliency Manager	The component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console. The Resiliency Manager is deployed as a virtual appliance.
Infrastructure Management Server (IMS)	The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance.
	To achieve scale, multiple IMSs can be deployed in the same data center.
Veritas InfoScale Operations Manager Management Server	The component that allows discovery of InfoScale applications that are already configured in Veritas InfoScale Operations Manager. Also referred to as Veritas InfoScale Operations Manager server.
	You can manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager on Linux, Solaris, AIX as well as Windows platform.
	See "Managing InfoScale applications using Resiliency Platform" on page 27.

Replication Gateway	The component of Veritas Resiliency Platform Data Mover that is deployed as a virtual appliance on both data centers and used to perform replication between the data centers.
data center	For a disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.
asset infrastructure	The data center assets that you add to Resiliency Platform for discovery and monitoring by the IMS.
	The asset infrastructure can include hosts (Windows or Linux servers), virtualization servers for Hyper-V and VMware, and enclosures (storage arrays). Once the asset infrastructure is discovered by the IMS, the discovered virtual machines or applications are listed in the console as assets to manage or protect.
resiliency group	The unit of management and control in Resiliency Platform. You organize related assets into a resiliency group and manage and monitor them as a single entity.
service objective	A template to define the type of operations and technologies that are supported for a group of assets. You apply a service objective to each resiliency group.
	A template which identifies the characteristics of a service. These could be availability related characteristics such as local redundancy, and number of nodes in a cluster or DR characteristics such as remote recovery, Recovery Point Objective (RPO) SLAs, rehearsal support etc. Service objective is applied when a group of assets are being added to a resiliency group.
	Resiliency Platform monitors the resiliency groups based on the service objective definition and raises the risks as applicable.
Virtual Business Service (VBS)	A multi-tier business service where each VBS tier hosts one or more resiliency groups. A VBS lets you group multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. VBS uses the vertical grouping mechanism to group the multiple services.You can also perform operations such as migrate, takeover, resync, rehearsal on the entire VBS.

About Resiliency Platform capabilities

Resiliency Platform helps you monitor and manage disaster recovery across multiple data centers. It provides the following capabilities.

Table 1-1   Resiliency Platform capabilities

Capability	More information
Configuring virtual machines and applications for remote recovery operations or basic monitoring	See "Managing applications for basic monitoring" on page 31. See "Managing applications for remote recovery (DR)" on page 57.
Starting and stopping resiliency groups for maintenance	See "Starting a resiliency group" on page 32. See "Stopping a resiliency group" on page 33.
Rehearsing disaster recovery	See "Performing the rehearsal operation" on page 61.
Migrating a resiliency group	See "Migrating a resiliency group of applications" on page 63.
Taking over resiliency groups	See "Taking over a resiliency group of applications" on page 64.
Performing the resync operation	See "Performing the resync operation" on page 64.
Managing activities and resiliency plans	See "Managing activities" on page 79. See "Managing resiliency plans" on page 81.
Displaying an overview of your resiliency domain including the number and health of your resiliency groups	See "About the Resiliency Platform Dashboard" on page 66. See "Displaying resiliency group information and status" on page 34.
Monitoring risks for protected assets	See "About risk insight" on page 70.
Viewing reports	See "Viewing reports" on page 69.

About permissions for operations in the console

Users that are configured for Resiliency Platform have permission by default to view the web console but not to perform any operations. Permissions for operations must be assigned separately by a Resiliency Platform administrator, who assigns

the appropriate personas to users or groups. A persona is a role with access to a set of operations. The administrator can further limit the scope of some operations by selecting the objects, such as resiliency groups, to which the user has access.

For example, an administrator can assign one user the permission to perform operations on resiliency group RG1 and assign another user the permission to perform operations on RG2. If more resiliency groups are added later, the administrator needs to update permissions to assign access to the new resiliency groups.

Some objects, such as resiliency plans or virtual business services, can include multiple resiliency groups. To perform an operation on such an object, a user must have access to all its resiliency groups. Otherwise, the operation fails.

For more information on setting up user access to operations, refer to the *Deployment Guide*.

Managing applications using Resiliency Platform

This chapter includes the following topics:

- Managing applications using Resiliency Platform

- Providing inputs for partially discovered applications

- Managing custom applications

- About application bundles

- Adding an application bundle to the Resiliency Manager

- Removing an application bundle

- Installing an application bundle on selected application hosts

- Uninstalling an application bundle from selected hosts

- Enabling and disabling application bundle on selected application hosts

- Managing auto-deploy for an application bundle

- Editing the discovery schedule for an application type

- Viewing the details of application types

- Viewing the applicable host details

Managing applications using Resiliency Platform

You can use the Veritas Resiliency Platform to manage and protect your applications that are configured in the resiliency domain. For more information on supported

applications and their versions refer to *Hardware and Software Compatibility List (HSCL)*. The Resiliency Platform supports application discovery on physical hosts as well as VMware and Hyper-V virtual machines.

When hosts are added to and discovered by the Infrastructure Management Server (IMS), applications residing on those hosts are displayed on the Resiliency Platform. They are listed on the **Unmanaged** tab. Note that for discovering Oracle instances, the `oratab` file must be present and must contain the entries for oracle applications.

For certain application instances, you need to provide additional information such as application database file path, or user name and password to complete the discovery. These are listed on the **Application** tab with a **Pending Inputs** warning.

See "Providing inputs for partially discovered applications" on page 16.

Resiliency Platform has inbuilt support for Microsoft SQL Server (MSSQL) and Oracle applications. To discover, manage, and protect other applications, you need to add them using the **Add Custom Application** wizard.

See "Managing custom applications" on page 16.

You can also discover and manage applications using the Application Enablement Software Development Kit (SDK). You can create a bundle with predefined Perl APIs and upload the same on the Resiliency Manager. You can either auto deploy the bundle on all the associated hosts in your data center or do it manually by selecting the managed hosts. The Perl scripts in the bundle developed using the SDK discover and report the applications on the Resiliency Platform console. For more information about developing the bundle, refer to *Veritas Resiliency Platform 1.1: Application Enablement SDK* guide.

See "Adding an application bundle to the Resiliency Manager" on page 19.

Resiliency Platform lets you manage applications by grouping them into resiliency groups. Some examples of Resiliency Platform operations are create a resiliency group, edit the resiliency group to add or remove applications, start and stop the resiliency groups and so on. Applications must be completely discovered to add them into resiliency groups.

See "About resiliency groups" on page 29.

Resiliency Platform provides disaster recovery (DR) specific operations to protect your applications that are grouped into a resiliency group. For example you can configure disaster recovery for the resiliency group and also migrate the resiliency group to another data center.

See "Understanding asset types" on page 68.

See "Managing applications for basic monitoring" on page 31.

See "Managing applications for remote recovery (DR)" on page 57.

Providing inputs for partially discovered applications

When hosts are added to Resiliency Platform and discovered by the Infrastructure Management Server (IMS), applications residing on those hosts are displayed on the web console. But certain application instances are not completely discovered until you provide additional information such as application database file path, or user name and password.

Using the Resiliency Platform console, you can provide inputs to such partially discovered applications to enable complete discovery. Complete discovery of applications is essential to group them into a resiliency group and thereby perform the disaster recovery operations.

To provide inputs for partially discovered applications

1   Navigate

⚙   **Settings** (menu bar) > **Infrastructure** > **Details View**

Managed Host > **Application** tab

2   Applications with pending inputs are listed in the **Applications with inputs** table.

3   Right-click the partially discovered application and select **Enter Inputs**.

4   In the **Enter Inputs** panel, enter the required information, and click **Submit**.

See "Managing applications for basic monitoring" on page 31.

Managing custom applications

By default, Resiliency Platform discovers Microsoft SQL Server on Windows, Oracle applications on Linux, and applications configured in Microsoft Failover Cluster (Windows only). In certain circumstances, however, some Oracle applications may not be discovered. In addition, you may want to add other applications to Resiliency Platform. To manage and protect the applications that are not discovered by default, you need to add them using the **Add Custom Application** wizard.

Note: Adding custom applications using Microsoft Failover Cluster nodes is not recommended and not supported.

To add a custom application

1 Prerequisites

Do the following:

- For Linux, create a script to start, stop, and monitor the application. Resiliency Platform interacts with the application using these scripts. The scripts should reside on the same host as the application. For Windows applications, provide the application path such as `c:\windows\system32\notepad.exe`. A script is not required.

- Note the complete path to each script.

- Determine the user who should run the script. Often, it is the admin user or root user. Note the password of this user.

- Identify which data directory paths the application uses.

2 Navigate

⚙ **Settings** (menu bar) > **Infrastructure** > **Details View**

Managed Host > **Application** tab

3 In the **Application Hosts** table, select the host on which you want to add the custom application.

4 Click the vertical ellipsis and select **Add custom application**.

5 In the **Application type** drop-down, select the required application type.

6 On the **Fill inputs for application** page, do the following:

- Verify that you select the correct data center and host.

- Use the information you collected in step 1 to complete the form.

- Specify the instance name.

Note that if you selected Oracle as **Application type** in step 5, you need to specify same name (SID) in both the Instance name fields. If you specify different names for both the instance name fields, your application may not be discovered.

7 Click **Submit**.

After you add a custom application, you organize it with other applications into a resiliency group.

To edit a custom application

1   Navigate

⚙   **Settings** (menu bar) > **Infrastructure** > **Details View**

Managed Host > **Application** tab

2   In the **Custom Applications** table, select the application that you want to edit.

3   Click the vertical ellipsis and select **Edit Application**.

4   In the **Edit custom application inputs** panel, make the required changes and click **Next**.

5   Click **Finish**.

To delete a custom application

1   Navigate

⚙   **Settings** (menu bar) > **Infrastructure** > **Details View**

Managed Host > **Application** tab

2   In the **Custom Applications** table, select the application that you want to delete.

3   Click the vertical ellipsis and select **Delete Application**.

4   In the **Confirm delete custom application** review the information and click **Next**.

5   Click **Finish**.

See

See

About application bundles

Using the Resiliency Platform console you can manage the discovered applications that are installed in your data center.

For some applications such as MS SQL and Oracle the discovery scripts are pre-bundled with the Resiliency Platform. For others, you can create a bundle using the Application Enablement SDK to discover and manage the applications. A sample

bundle is available for downloading on the **Application Types** page in **Settings**>**Application Support**.

The pre-bundled applications are listed on the **Pre-bundled** tab whereas the bundles that are created using Application Enablement SDK and are uploaded on Resiliency Manager are listed on the **Uploaded** tab.

Pre-bundled applications are installed on the applicable hosts whenever any application host is added to the Resiliency Manager. These applications cannot be added or removed later but can be installed, uninstalled, enabled, or disabled on selected application hosts.

The applications that are created using Application Enablement SDK can be added and removed. You can install, uninstall, enable, or disable them on selected application hosts.

The auto deploy feature lets you install an application automatically on every new host that you add to your data center. You can enable or disable this feature for a particular application. This feature is available only for the applications created using Application Enablement SDK. For pre-bundled applications, the auto-deploy option is set to Yes by default.

Adding an application bundle to the Resiliency Manager

Using Application Enablement SDK you can create an application bundle to discover and manage applications in your data center. You can upload this bundle on the Resiliency Manager and install it on all the applicable hosts that are associated with the platform. You can also choose to install the bundle on selected application hosts.

The bundle should be in `.tar.gz` format.

For more information on how to create the bundle, refer to *Veritas Resiliency Platform 2.0: Application Enablement SDK* guide.

After you upload the bundle, you can view the following information on the **Uploaded** tab on **Application Types** page:

- Application type
- Application category
- Version
- Platform
- Whether the application bundle is enabled for auto deploy.

- The Status column displays the installed versus applicable hosts data.

To add an application bundle to the Resiliency Manager

1 Navigate

    ⚙   **Settings** (menu bar)

    Under **Infrastructure Settings**, click **Application Support**.

2 On the **Uploaded** tab, click **Add**.

3 In the **Add Application Type** panel, click **Browse** to select the application bundle. You can upload only one file at a time. Click **Upload**.

4 Review the application information in the table.

5 Do the following and click **Submit**.

- You can choose to auto deploy the applications on all the associated hosts.

- If an application type with the same version already exists, then you can choose to overwrite it. If the lower version of the application exists, you can choose to upgrade to the higher version.

See "Installing an application bundle on selected application hosts" on page 21.

Removing an application bundle

Using the Resiliency Platform console, you can remove the application bundle that you had created Application Enablement SDK from the Resiliency Manager.

You cannot remove an application bundle that is installed if the managed host is disconnected from the Infrastructure Management Server (IMS).

To remove the application bundle

1 Navigate

    ⚙   **Settings** (menu bar)

    Under **Infrastructure Settings**, click **Application Support**.

2 On the **Uploaded** tab, select the application which you want to remove, and click **Remove**.

3 In the **Remove Application Type** panel, review your selection and click **Submit**.

Installing an application bundle on selected application hosts

Using the Resiliency Platform console, you can install an application bundle on selected application hosts. The application bundle can be pre-bundled or created using Application Enablement SDK.

To install an application bundle on selected application hosts

1   Navigate

   ⚙   **Settings** (menu bar)

      Under **Infrastructure Settings**, click **Application Support**.

2   On the **Uploaded** or the **Pre-bundled** tab, double-click on the application which you want to install on selected application hosts.

3   On the application type details page, select the hosts on which you want to install the bundle and click **Install**.

   Only those hosts which match the platform that is defined for the application are listed in the table.

4   On the **Install Application Type** panel, review your selection, and click **Submit**.

Uninstalling an application bundle from selected hosts

Using the Resiliency Platform console, you can uninstall the application bundle from the selected hosts. The application bundle can be pre-bundled or created using Application Enablement SDK.

To uninstall an application bundle from selected application hosts

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

2 On the **Uploaded** or the **Pre-bundled** tab, double-click on the application which you want to uninstall.

3 On the application type details page, select the hosts from which you want to uninstall the bundle and click **Uninstall**.

4 On the **Uninstall Application Type** panel, review your selection, and click **Submit**.

See

Enabling and disabling application bundle on selected application hosts

Using the Resiliency Platform console, you can enable and disable the application bundle on the selected hosts to resume or pause the discovery of the application instances.

When you disable the application bundle, the application discovery is paused. The current state of the application is not displayed on the Resiliency Platform console.

To enable or disable application bundle on selected application hosts

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

2 On the **Uploaded** or the **Pre-bundled** tab, double-click on the application type which you want to enable or disable.

3 On the application details page, select the hosts on which you want to enable or disable the bundle and click the appropriate menu option.

■ **Enable**

- **Disable**

4   On the **Application Type** panel, review your host selection, and click **Submit**.

See "Installing an application bundle on selected application hosts" on page 21.

Managing auto-deploy for an application bundle

The auto deploy feature lets you install the application on every new host that you add to your data center. You can enable or disable this feature for a particular application. This feature is available only for the applications created using Application Enablement SDK. For pre-bundled applications, the auto-deploy option is set to Yes by default.

To manage auto-deploy for an application bundle

1   Navigate

⚙   **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

2   On the **Application Types** page, on the **Uploaded** tab, right-click the application type and select **Set/Unset Auto-Deploy**.

3   On the **Set/Unset Auto-Deploy** page click **Submit**.

See "Enabling and disabling application bundle on selected application hosts" on page 22.

See "Editing the discovery schedule for an application type" on page 23.

Editing the discovery schedule for an application type

Using the Resiliency Platform console, you can edit the discovery schedule for an application type. Application discovery is of two types: deep and probe. Deep discovers the entire application and its components including files. Probe only checks the status of the application instances. For example whether the application is online or offline.

The discovery frequency for deep can range from 240 minutes to 1440 minutes. For probe, the range is one minute to 60 minutes.

The default discovery frequencies for deep and probe are 360 and 10 minutes respectively.

To edit the discovery schedule for an application type

1   Navigate

⚙   **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

2   On the **Application Types** page, on the **Uploaded** tab, right-click the application type and select **Properties**.

3   On the **Properties** page, click on the **Frequency** column to change the discovery schedule for Deep and Probe, and click **Submit**.

See "Installing an application bundle on selected application hosts" on page 21.

Viewing the details of application types

Using the Resiliency Platform console you can view the details of the application types.

You can view the following information on the **Uploaded** tab:

- Application type

- Category such as database

- Version

- Platform information

- Whether the application bundle is enabled or disabled for auto deploy on all the applicable hosts.

- The **Status** column displays the installed versus applicable hosts data.

You can view the following information on the **Pre-bundled** tab:

- Application type

- Category such as database

- Whether the application is enabled or disabled for auto deploy on all the applicable hosts.

To view the details of application types

◆   Navigate

⚙   **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

See "Adding an application bundle to the Resiliency Manager" on page 19.

See "Installing an application bundle on selected application hosts" on page 21.

See "Managing auto-deploy for an application bundle" on page 23.

Viewing the applicable host details

Using the Resiliency Platform web console you can view the details of the hosts on which an application bundle can be installed.

You can view the following information of the hosts in this view:

■   Host name

■   Platform

■   Family

■   Architecture

■   Status of the bundle on the host. Whether it is enabled, disabled, or not installed.

To view the applicable host details

1   Navigate

⚙   **Settings** (menu bar)

Under **Infrastructure Settings**, click **Application Support**.

2   On the **Application Type** view, double-click on an application.

See "Adding an application bundle to the Resiliency Manager" on page 19.

See "Installing an application bundle on selected application hosts" on page 21.

See "Managing auto-deploy for an application bundle" on page 23.

Managing InfoScale applications using Resiliency Platform

This chapter includes the following topics:

- About Veritas InfoScale Operations Manager

- Resiliency Platform support for InfoScale applications

- Managing InfoScale applications using Resiliency Platform

About Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager gives you a single, centralized management console for the Veritas InfoScale products. You can use it to monitor, visualize, and manage storage and cluster resources, and generate reports about these components in the Management Server domain. Veritas InfoScale Operations Manager helps administrators centrally manage diverse data center environments.

Resiliency Platform support for InfoScale applications

A typical workflow of Veritas Resiliency Platform for InfoScale applications consists of a Veritas InfoScale Operation Manager server reporting to a Resiliency Manager. The InfoScale applications should be already configured in Veritas InfoScale Operations Management server. You can group the InfoScale applications into resiliency groups or VBSs to recover, monitor, visualize, and generate reports about these applications in the Resiliency Platform.

Note: Only the applications that are managed by InfoScale Availability (VCS) are supported in Veritas Resiliency Platform.

The following diagram depicts the general workflow of configuring the InfoScale applications using Resiliency Platform.

Figure 3-1   A typical workflow for recovering managed InfoScale applications
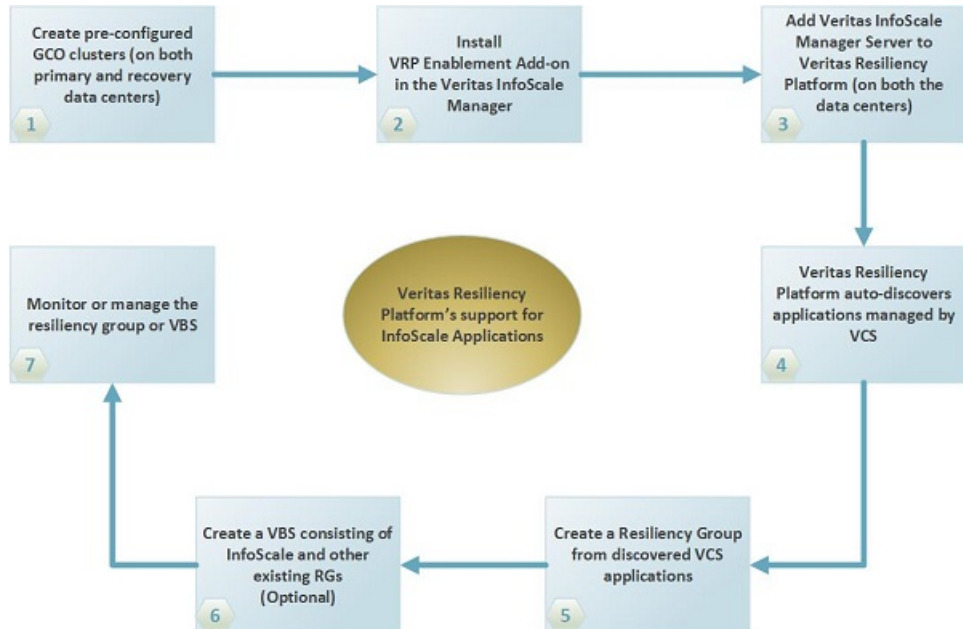
Managing InfoScale applications using Resiliency Platform

Veritas Resiliency Platform lets you manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager Management Server. The InfoScale applications are automatically discovered in the Resiliency Platform when the Veritas InfoScale Operations Manager server is added to the resiliency domain. They are listed on the **Assets** page under the **Unmanaged** tab. You can filter the InfoScale applications using the **InfoScale applications** check box under the **More Options** drop-down menu.

For more information on adding a Veritas InfoScale Operations Manager server to the Resiliency Domain, refer to *Veritas Resiliency Platform Deployment Guide*.

Veritas InfoScale Operations Manager users must download and install Veritas Resiliency Platform Enablement add-on to automatically discover the InfoScale applications. You can download the add-on from Veritas Services and Operations Readiness Tools (SORT). You cannot add or modify InfoScale applications through Resiliency Platform. They can be added or modified only by an administrator through Veritas InfoScale Operations Manager.

Managing resiliency groups

This chapter includes the following topics:

About resiliency groups

Resiliency groups are the unit of management and control in Veritas Resiliency Platform. After assets are added to Resiliency Platform, you organize related assets into a resiliency group that you can protect and manage as a single entity.

For example, you can organize several applications into a resiliency group and name it `SQL_Server_Group`. Then, when you perform an operation on `SQL_Server_Group` from the Resiliency Platform console, all the applications in the group are affected. For example, if you start `SQL_Server_Group`, all the applications

in the group start. Similarly, you can organize virtual machines into a resiliency group and perform operations that affect all the virtual machines in the group.

Note: A resiliency group must contain similar types of objects, either all applications or all virtual machines. It cannot contain a mix of the two.

The operations available for a resiliency group depend on how it is configured. During the configuration of a resiliency group, you apply a service objective that identifies the objective or intent for that group of assets. If you apply a service objective that supports remote recovery, the resiliency group supports operations like migrate and take over.

You can optionally use a service objective that only monitors the assets and provides only basic operation capabilities like start and stop operations and no remote recovery operations.

See "About service objectives" on page 31.

See "Managing applications for basic monitoring" on page 31.

See "Managing applications for remote recovery (DR)" on page 57.

Prerequisites for creating resiliency groups with applications

Following prerequisites are for creating resiliency groups with applications:

- Applications must be completely discovered.

- For Microsoft SQL Server, the **Guest** user must have **Connect** permission on all the databases, to create a resiliency group.

Prerequisites for creating resiliency groups with InfoScale applications

Following prerequisites are for creating resiliency groups with InfoScale applications:

- You cannot add both InfoScale and non-InfoScale applications within a same resiliency group.

- The selected InfoScale applications in the resiliency group should be from the same Global Cluster Option (GCO) pair. You cannot add InfoScale applications from different GCO pairs.

- The InfoScale applications in the resiliency group should be from different data centers.

About service objectives

Service objectives define the type of protection to be applied to a group of data center assets. For example, an option for remote recovery which allows assets being managed by a resiliency group to be recovered at a remove location (DR) using service objective can include operations such as migrate or take over.

Such service objectives can include tunables such as Recovery Point Objective (RPO) for assets being managed in that resiliency group.

When you create a resiliency group of assets in Veritas Resiliency Platform, you select a service objective to apply to that group of assets. The wizard then prompts you for any additional information that is needed to prepare the resiliency group for the supported operations.

For example, if you apply a recovery service objective that supports recovery of virtual machines to a remote data center, you must select the recovery data center.

Service objectives are provided as templates that must be activated before use. A set of pre-activated service objectives with default settings is provided.

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

Note: Authorization to perform operations also depends on the permissions assigned to users and groups in Resiliency Platform.

For more information on customizing service objectives, refer to the Deployment Guide.

Managing applications for basic monitoring

When you create a resiliency group, you select a service objective that specifies the operations supported for that resiliency group.

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

This topic explains how to configure a resiliency group for basic monitoring.

Configuring a resiliency group for remote recovery has additional prerequisites and steps and is described in a separate topic.

See "Managing applications for remote recovery (DR)" on page 57.

To manage applications for basic monitoring

1   Prerequisites

   ■ The asset infrastructure must be added to Resiliency Platform and discovery of the applications must be complete. The applications should be running. For more information on adding asset infrastructure, refer to the *Deployment Guide*.

   ■ For InfoScale applications, the asset infrastructure must be added to the Veritas InfoScale Operations Manager and the discovery of the applications must be complete. For more information, refer to the Veritas InfoScale Operations Manager documentation.

   ■ For InfoScale applications, Veritas Cluster Servers (6.0 onwards) with Global Cluster Option (GCO) should be enabled.

2   Navigate

   ▭   **Assets** (navigation pane) > **Resiliency Groups** tab > **Manage & Monitor Virtual Machines or Applications**

   You can also launch the wizard from the **Unmanaged** or **Overview** tabs.

3   Select the applications:

   ■ Select **Applications** as the asset type, select the data center, and select other filters as needed to display a list of applications.

   ■ Drag and drop the applications to **Selected Instances**.

4   Select the service objective that provides monitoring and start and stop operations only.

5   Supply a name for the resiliency group and submit.

6   Verify that the new resiliency group is added to the **Resiliency Groups** tab.

   Optionally, use **Recent Activities** (bottom pane) > **Details** to view the details of this task in a graphical representation.

Starting a resiliency group

When you start a resiliency group, you start all the underlying assets in it.

To start a resiliency group

1   Navigate

⊟   **Assets** (navigation pane) > **Resiliency Groups** tab

2   Locate your resiliency group. Use filters or Search as needed.

3   ⋮   On the row for the resiliency group, select the vertical ellipsis > **Start**.

You can also perform operations from the Details page.

4   On the **Start Resiliency Group** wizard, select the data center in which to start the group, and submit.

If you have applied update 2.0.0.100 on Veritas Resiliency Platform 2.0, you can select the checkbox on the **Start Resiliency Group** wizard to start the post-replication operations of migrate or takeover workflow on the production data center such as refreshing storage, network, compute, and customization.

To display a record and a graphic representation of what you did, select the **Recent Activities** at the bottom of the page, find your task, and select **Details**.

5   If necessary, notify users after you start the resiliency group.

Stopping a resiliency group

When you stop a resiliency group, you stop all the assets that make up the group.

A typical reason for stopping a resiliency group would be to update or perform maintenance in one or more of the underlying assets.

To stop a resiliency group

1   Prerequisites

■   Make sure that you are aware of all the assets in the resiliency group, and the potential effect on users if you shut them down.

■   Choose a time for stopping the resiliency group that minimizes any disruption of service.

■   If necessary, notify users before you stop the resiliency group.

2   Navigate

⊟   **Assets** (navigation pane) > **Resiliency Groups** tab

3 Locate the resiliency group. Use filters or Search as needed.

4 ⋮ On the row for the resiliency group, select the vertical ellipsis > **Stop**.

You can also perform operations from the Details page.

5 On the **Stop Resiliency Group** screen, select the data center in which to stop the resiliency group, and submit.

To display a record and a graphic representation of what you did, select the **Recent Activities** at the bottom of the page, find your task, and select **Details**.

Displaying resiliency group information and status

You can display resiliency group information and status in the following ways:

Table 4-1 Displaying resiliency group information and status

Location	Level of detail	Useful for
Resiliency Platform Dashboard	Lowest. Displays the number of resiliency groups under Resiliency Platform control and the total number of groups in error, at risk, and healthy.	Getting a quick overview of the resiliency group population and health throughout Resiliency Platform. See "About the Resiliency Platform Dashboard" on page 66.
▢ **Assets** > **Resiliency Groups** tab	Medium. Lists all your resiliency groups in one place.	Seeing what is in each of your data centers, the state of the groups, and so on.

Table 4-1   Displaying resiliency group information and status *(continued)*

Location	Level of detail	Useful for
Resiliency group-specific screen	Highest. Lists each asset in the resiliency group, their type, and state.	Getting detailed information on a resiliency group and its underlying assets, including disaster recovery status. This screen lists available operations for the group. See "Viewing resiliency group details" on page 38.

This section discusses the second method of displaying resiliency group information and status: using the **Assets** page. The **Assets** page gives you a quick overview of all your resiliency groups.

To display resiliency group information and status

1 Navigate

⊞ **Assets** (navigation pane) > **Resiliency Groups** tab

2 Review information and status

For a quick health check of your resiliency groups, review the colored boxes above the table. Select a box to show only the resiliency groups in that category; for example, select the green square to display only the resiliency groups that are healthy.

Blue	The total number of resiliency groups
Yellow	The number of resiliency groups at risk
Green	The number of resiliency groups that are healthy

By default, the table lists all resiliency groups. Use the drop-list and search field to filter your results, and select a table heading to sort the groups.

In the table, the key fields are **State**, **Service Objective**, and **Data Availability**. Possible states are:

Status	**Normal** - the assets within the resiliency group are normal.
	At Risk - the assets within the resiliency group are at risk.
State	**Online** - The assets within the resiliency group are running.
	Partial - One or more of the assets in the resiliency group are offline.
	Offline - The assets in the resiliency group are powered off or not running.
Active DC	Name of the active data center.
Type	Application Group: The resiliency group comprises of applications.
	Virtual Machine Group: The resiliency group comprises of virtual machines.
Service Objective	Service objective selected for the resiliency group.
Data Availability	Resiliency Platform supports several replication technologies.
	If no replication type is shown, consider configuring replication.

Viewing InfoScale applications details

Veritas Resiliency Platform lets you manage the InfoScale applications that are added in Veritas InfoScale Operations Manager (7.0 onwards).

To view InfoScale applications information

1   Navigate

   Assets > **Unmanaged** tab

2   Use one or more of the following drop-downs to filter your list of applications:

Asset Type	Select application.
Data center	Select the data center in which the application is located.
Application Type	Select the application type.
	In InfoScale applications, Application Type resources with service group dependencies are considered as a single application.
	Note: Only the service group with resources that are of Application type are discovered under the Unmanaged tab.

3   Under **More Options** drop-down menu, select **InfoScale applications**.

4   At the end of the InfoScale application row,

⋮   Click the vertical ellipsis, and select **More Information**.

5   In the **More Information** window, you can view the following InfoScale details:

Details	It displays the Application Name along with its Host Name, and State (Offline or Online). It also provides the type of the application.
Availability	It displays the type of availability product used, name, and type of the service group. It also provides the system list, cluster details of the application and fire drill service related information.
Data protection	It displays detailed information about the replication details, such as Replication Type, Replication availability, of the application.

Viewing resiliency group details

Using the Resiliency Platform console, you can view detailed information on each of your resiliency groups. The overall health of the resiliency group, its underlying assets and their current state is displayed.

Resiliency group for which disaster recovery (DR) operation is configured successfully, you can view information which includes the state of the replication for the resiliency group (for example, synchronized), used replication technology (for example, EMC SRDF), associated alerts, the details about the applications or the virtual machines in the resiliency group, replication lag, recovery time, and so on.

Note that for virtual machines, the recovery time is available only after the rehearse operation is complete.

To view details of a resiliency group

1 Navigate

⊟ **Assets** (navigation pane) > **Resiliency Groups** tab

2 Locate your resiliency group. Use filters and search as needed.

3 On the row for the resiliency group, select the vertical ellipsis > **Details**.

You can also double-click the row to view details.

The details page includes the following:

- Menu options for operations that you can perform on the resiliency group.
- Details of how the resiliency group is configured.
- Status information.
- A list of the resiliency group assets and their state.

See

Editing a resiliency group

You can edit the resiliency group information including the group name as well as change the underlying assets on which the resiliency group is based when the resiliency group is configured for basic monitoring using the Monitor assets service objective.

If the resiliency group is already protected for DR, then the wizard proceeds with the DR configuration letting you make any changes if required.

To edit the resiliency group information

1 Prerequisites

- Determine the potential impact editing the resiliency group may have on users. If necessary, notify users of the upcoming change.

2 Navigate

⊟ **Assets** (navigation pane) > **Resiliency Groups** tab

3   Locate the resiliency group. Use filters or Search as needed.

4   On the row for the resiliency group, select the vertical ellipsis > **Edit**.

You can also edit the resiliency group from its Details page.

The steps for editing the resiliency group are the same as creating it.

Deleting a resiliency group

When you delete a resiliency group from Resiliency Platform management, you can no longer monitor, manage, or protect it using Resiliency Platform. Deleting the resiliency group from Resiliency Platform has no effect on the underlying assets.

To successfully complete the delete operation ensure the following:

- The assets on the production data center are running and accessible.

- The xprtld daemon on the virtual machines is running.

On successful completion of the delete operation, you will notice the following:

- During the operation, replication is stopped and Veritas Replication Sets are deleted on gateways and on-premises virtual machines.

- Journal disks are removed from the virtual machines on the production data center and cloud virtual machines instances are deleted. For physical machines, the journal disks are unsigned.

- All the cloud virtual machines disks that are attached to the cloud Replication Gateway are deleted.

Note: Replication Gateway pairs are not deleted during the delete operation. If required you can delete the pair from the **Gateway Pair** details page.

To delete a resiliency group

1   Navigate

Assets (navigation pane) > **Resiliency Groups** tab

2   Locate the resiliency group. Use filters or Search as needed.

3   On the row for the resiliency group, select the vertical ellipsis > **Delete**.

You can also perform operations from the Details page

4   Confirm the deletion.

Preparing for disaster recovery configuration

This chapter includes the following topics:

- An overview of key steps required for disaster recovery of applications

- Prerequisites for configuring applications for disaster recovery

- An overview of key steps required for disaster recovery of InfoScale applications

- Prerequisites for configuring InfoScale applications for disaster recovery

- About replication technologies used in disaster recovery of applications

- Configuring DNS server settings for a data center

An overview of key steps required for disaster recovery of applications

This section lists the key steps required to configure the disaster recovery for applications using Resiliency Platform.

Table 5-1   Disaster recovery for applications - an overview of key steps

Action	Description	Refer to
Set up your replication environment	Configuration before you install Resiliency Platform. It includes configuring applications on physical and virtual machines, configuring replication and so on.	See "Configuring application disaster recovery using EMC SRDF replication" on page 47. See "Configuring application disaster recovery using NetApp SnapMirror replication" on page 51. See "Configuring application disaster recovery using EMC RecoverPoint replication" on page 52.
Add the asset infrastructure	Add the asset infrastructure to the Infrastructure Management Server (IMS) using the Resiliency Platform web console.	Refer to the *Deployment Guide* .
Configure your assets for disaster recovery	Group the required applications in a resiliency group and enable disaster recovery for the resiliency group.	See "Managing applications for basic monitoring" on page 31. See " Editing a resiliency group" on page 39. See "Managing applications for remote recovery (DR)" on page 57.
DR operations	Perform the required DR operations: Migrate and takeover.	See "Migrating a resiliency group of applications" on page 63. See "Taking over a resiliency group of applications" on page 64.

Prerequisites for configuring applications for disaster recovery

To be able to perform disaster recovery (DR) operations on the applications in the data center, ensure that the following requirements are met.

- The application host must be added to Resiliency Platform and discovery of the applications must be complete. The applications should be running.

- Add the discover hosts and the enclosures for replication to the IMS at the production and recovery data center.

- If applications are installed on VMware virtual machines, you need to configure the VMware vCenter for discovery. Similarly if the applications are installed on Hyper-V virtual machines, you need to configure the Hyper-V server for discovery.

- The replicated storage provided to the application hosts from the storage arrays is discovered if you have configured the discovery hosts and the enclosures. Replicated storage can be provisioned to the application in numerous ways. For applications installed inside virtual machines, Resiliency Platform discovers storage that is provisioned to the virtual machines in raw mode (e.g. VMware RDM mode) from the hypervisors. Raw mode can also be achieved by using NPIV HBAs and provisioning storage (LUNs) directly from the arrays to the Virtual HBA Ports assigned to the virtual machines.

  If the storage is provisioned to the application virtual machines using hypervisor, then you need to provision appropriate snapshot devices to the virtual machines in the same manner as the replicated devices. This is required to perform the rehearsal operation successfully.

- For NetApp SnapMirror and HP 3PAR RemoteCopy, the snapshot devices do not exist beforehand and are created as a part of the Rehearsal operation. These snapshot devices should be mapped to the application hosts. As part of the Rehearsal operation, Resiliency Platform maps the snapshot devices to the same hosts to which the replicated LUNs are mapped. If the replicated LUNs are mapped through the hypervisor, then Resiliency Platform cannot further provision them to the application virtual machines.

  As a result, Rehearsal is not supported with these replication technologies when the application is installed in virtual machines with replicated storage mapped as raw LUN by the hypervisor.

- Applications must be members of the same Veritas Replication Set.

- They must use the same replication technology.

- Application binaries should be stored on local storage and data files on replicated storage.

- The DNS server settings should be configured for both data centers. DNS settings are required for binding a host name to different IP addresses on the DR site. This is required only if you plan to use the Resiliency Platform for performing DNS updates.

- For applications on Windows with Failover Clustering, you may have to plumb or unplumb the IP addresses using appropriate Failover Clustering roles. If Failover Clustering is not used, you must plumb the application IP addresses on all the systems across the data centers. The Resiliency Platform console does not manage plumbing or unplumbing of IP address for applications.

- Ensure that you have disabled the 'Quick removal' policy for disks in Windows Server 2008 R2 and disabled the 'write-cache' policy for disks in Windows Server 2012 R2.

- Applications should be preconfigured at the recovery data center before creating resiliency groups and proceeding with the DR operations. Discovery of these applications on the recovery data center must be complete and status should be offline. Ensure that the application instance name should be same at the production and recovery data center.

- Linux LVM Volume Group used for data file storage of an application should not be in exported state.

- On Linux application hosts, the data mount point of the application should be present in /etc/fstab file.

See "About resiliency groups" on page 29.

See "Managing applications for remote recovery (DR)" on page 57.

See "About replication technologies used in disaster recovery of applications" on page 47.

An overview of key steps required for disaster recovery of InfoScale applications

This section lists the key steps required to configure the disaster recovery of InfoScale applications using Resiliency Platform.

Table 5-2   Disaster recovery for InfoScale applications - an overview of key steps

Action	Description	Refer to
Ensure the prerequisites	Create pre-configured GCO clusters on the production and the recovery sites	Refer to the *Veritas™ Cluster Server Administrator's Guide*.

Table 5-2   Disaster recovery for InfoScale applications - an overview of key steps *(continued)*

Action	Description	Refer to
Add the asset infrastructure	Install the Veritas Resiliency Platform Enablement add-on add-on on Veritas Infoscale Operations Manager server Add the Veritas Infoscale Operations Manager servers on production and recovery site to the respective Resiliency Managers.	Refer to the *Deployment Guide* .
Configure your assets for disaster recovery	Group the applications discovered by Infoscale Operations Manager in a resiliency group and enable disaster recovery for the resiliency group.	See "Managing applications for basic monitoring" on page 31. See " Editing a resiliency group" on page 39. See "Managing applications for remote recovery (DR)" on page 57.
DR operations	Perform the required DR operations: Migrate and takeover.	See "Migrating a resiliency group of applications" on page 63. See "Taking over a resiliency group of applications" on page 64.

Prerequisites for configuring InfoScale applications for disaster recovery

To be able to perform disaster recovery (DR) operations on the InfoScale applications in the data center, ensure that the following requirements are met.

- For InfoScale applications, VCS Clusters must be added to Veritas InfoScale Operations Manager on both the production and recovery data centers. The discovery of the applications must be complete. Veritas Resiliency Platform Enablement add-on must be installed on all the clusters that are managed by Resiliency Platform on the production and the recovery data centers. For more information, refer to the Veritas InfoScale Operations Manager documentation.

■ For managing recovery of InfoScale applications, Veritas Cluster Servers (6.0 onwards) with Global Cluster Option (GCO) should be enabled.

About replication technologies used in disaster recovery of applications

For a successful disaster recovery operation of applications, you need to ensure that the data is synchronized between the primary and the secondary data centers. This is achieved using data replication.

The following table lists the supported configurations based on array replication and clustering:

Table 5-3   Supported configurations using EMC Symmetrix Remote Data Facility (SRDF), NetApp Snap Mirror, Hitachi True Copy, EMC RecoverPoint, HPE 3PAR Remote Copy replication

Operating System	Type of host	Clustering	Supported
Windows	Hyper-V virtual machine	Microsoft Failover Custering (MS FoC)	No
Windows	VMware virtual machine	MS FoC	No
Windows	Physical systems	MS FoC	Yes
Windows	Hyper-V virtual machine	Non - MS FoC	Yes
Windows	VMware virtual machine	Non - MS FoC	Yes
Windows	Physical systems	Non - MS FoC	Yes
Linux	Hyper-V virtual machine	NA	No
Linux	VMware virtual machine	NA	Yes
Linux	Physical systems	NA	Yes

Configuring application disaster recovery using EMC SRDF replication

This appendix includes the following scenarios:

- See the section called "Configuring application disaster recovery using EMC SRDF with Microsoft Failover Clustering" on page 48.

- See the section called "Configuring application disaster recovery using EMC SRDF without Microsoft Failover Clustering" on page 49.

Configuring application disaster recovery using EMC SRDF with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC SRDF when the hosts are a part of a Microsoft failover cluster. For EMC SRDF-based replication, all applications consuming storage from a consistency group must belong to the same resiliency group. A consistency group is a collection of Symmetrix LUNs that helps in maintaining write consistency during replication.

- Ensure that EMC Symmetrix Solutions Enabler (version v7.4, or later) is installed on a host and the SRDF device groups are already set up for the replication between the primary and remote arrays.

- Ensure that EMC Symmetrix Gatekeeper device is present on the array control host. You can designate any host including the Hyper-V server as the array control host.

Note: The SRDF R1 and R2 LUNs must be on different hosts from different data centers.

- Ensure to enable the Failover Cluster roles on the Windows Server 2012R2 hosts at the production and recovery data centers.

- Ensure that you have created Microsoft failover cluster using the required nodes at the production and recovery data centers.

- Ensure that you have created the cluster shared volume (CSV) on the replicated shared disk (R1) on the application server at the production data center. On the application hosts configured at the recovery data center, re-scan the storage on all the Microsoft failover cluster nodes.

- Configure application on the production data center's Microsoft failover cluster with their data on the replicated CSVs.

Once you have performed the necessary configurations, proceed with disaster recovery specific tasks.

Veritas Resiliency Platform configurations:

- Add Windows 2012 R2 hosts to the Infrastructure Management Server (IMS) using the **+ Add Application host** operation.

- Add the array control host where the SRDF device groups are configured, to the each IMS using the **+ Discovery Host** operation.

- Add Symmetrix enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI location on this discovery host. This operation returns the list of Symmetrix arrays (local and remote) accessible to the host. To configure disaster recovery for applications, select one or more local arrays.

  Default SymCLI location on Linux host   /opt/emc/SYMCLI/bin/

  Default SymCLI location on Windows host   C:\Program Files\EMC\SYMCLI\bin

  ---

  Note: Any physical or virtual machine can be designated as the array control host, including the virtual machine inside VMware ESX server that has EMC Symmetrix Gatekeeper device visibility.

- Perform add discover host and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitation: The rehearsal operation for resiliency groups using EMC Symmetrix Timefinder Snap is not supported in Microsoft Failover Cluster environment.

Configuring application disaster recovery using EMC SRDF without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC SRDF when the hosts are not a part of a Microsoft failover cluster.

- Ensure that EMC Symmetrix Solutions Enabler (version v7.4, or later) is installed on a host and the SRDF device groups are already set up for the replication between the primary and remote arrays.

- Ensure that EMC Symmetrix Gatekeeper device is present on the array control host. You can designate any host including the Hyper-V server as the array control host.

  Note: The replicated and primary LUNs must be on different hosts from different data centers.

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read and write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. Windows Storage Space Storage Pool is not supported.

- Ensure that you have configured application at the production data centre under the Hyper-V Manager and kept the data files on the replicated volumes.

- Ensure that the respective remote disks (Read only - R2 remote disk and snapshot) are in the offline state on the Hyper-V server at the recovery data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

Note: To perform the Rehearse operation, you must add the snapshot devices to the SRDF device group at the recovery data center, and thereafter map them to the application hosts at the recovery data center.

Once you have performed the necessary configurations, proceed with disaster recovery specific tasks.

Veritas Resiliency Platform configurations:

- Add the hosts to the Infrastructure Management Server (IMS) using the **+ Add Application host** operation.

- Add the host where the SRDF device groups are configured, to the Infrastructure Management Server (IMS) using the **+ Discovery Host** operation.

- Add Symmetrix enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI location on this discovery host. This operation returns the list of Symmetrix arrays (local and remote) accessible to the host. To configure disaster recovery for applications, select one or more local arrays.

  Default SymCLI location on Linux host   /opt/emc/SYMCLI/bin/

  Default SymCLI location on Windows host   C:\Program Files\EMC\SYMCLI\bin

---

Note: Any physical or virtual machine can be designated as the array control host, including the virtual machine inside VMware ESX server that has EMC Symmetrix Gatekeeper device visibility.

- Perform add discovery host and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitations

- EMC SRDF LUN-based replication, without device group, and replication in the adaptive copy mode are not supported.

- If the application hosts are inside the virtual machines, the replicated data disks must be mapped to these hosts in Raw mode only. Virtual disks are not supported.

- Logical grouping of disks, Windows Server Storage space storage pool, is not supported.

Configuring application disaster recovery using NetApp SnapMirror replication

This section lists the pre-requisites to enable the data replication using NetApp SnapMirror. For NetApp SnapMirror based replication, all applications that consume storage from a NetApp volume must belong to the same resiliency group.

- Ensure the NetApp volumes are already setup for replication between the primary and remote NetApp storage systems, and the replication has a replication schedule associated with it.
  Resiliency Platform does not support one NetApp volume having more than one SnapMirror destination volumes.

- Ensure to mount NetApp SnapMirror replicated volumes on the respective servers in both the sites. Do not mount the replicated peer NetApp Volumes on the same server.
  Also ensure that the volumes are replicated using SnapMirror policy type mirror. SnapMirror policy types vault and mirror-vault are not supported.
  NetApp share can be mounted with array IP or FQDN. For storage to application correlation to work successfully, ensure that the mount entry is consistent in the fstab.

- Change the permission to **Grant root access to all hosts** on the replicated volumes at both the production and recovery sites.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Resiliency Platform configurations:

- Ensure that all the application hosts are added and discovered completely by the Infrastructure Management Server (IMS).

- Add NetApp enclosure using the **+ NetApp Enclosure** option.
  Provide the discovery host name, NetApp storage system name or IP, and credentials.

- Ensure that all the application hosts are added and discovered completely by IMS at the disaster recovery (DR) site.

- Perform add enclosure operations for the IMS at the DR data center as well.

Configuring application disaster recovery using EMC RecoverPoint replication

This appendix includes the following scenarios:

- See the section called "Configuring application disaster recovery using EMC RecoverPoint with Microsoft Failover Clustering" on page 52.

- See the section called "Configuring application disaster recovery using EMC RecoverPoint without Microsoft Failover Clustering" on page 53.

Configuring application disaster recovery using EMC RecoverPoint with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC RecoverPoint when the hosts are a part of a Microsoft failover cluster. For RecoverPoint-based replication, all applications consuming storage from a consistency group must belong to the same resiliency group. A consistency group is a collection of LUNs that helps in maintaining write consistency during replication.

- Ensure that RecoverPoint consistency groups are set up on the control host for the replication between the primary and remote arrays.

- Ensure that EMC RecoverPoint groups are set up for the CRR replication between the primary and Secondary RecoverPoint Appliance.

- Ensure to enable the Failover Cluster roles on the Windows Server 2012R2 hosts at the production and recovery data centers.

- Ensure that you have created Microsoft failover cluster using the required nodes at the production and recovery data centers.

- Ensure that you have created the cluster shared volume (CSV) on the replicated shared disk on the application server at the production data center. On the application hosts configured at the recovery data center, re-scan the storage on all the Microsoft failover cluster nodes.

- Configure application on the production data center's Microsoft failover cluster with their data on the replicated CSVs.

Once you have performed the necessary configurations, proceed with disaster recovery specific tasks.

Veritas Resiliency Platform configurations:

- Add Windows 2012 R2 hosts to the Infrastructure Management Server (IMS) using the **+ Add application host** operation.

- Add the array control host where the RecoverPoint consistency groups are configured, to the each IMS using the **+ Discovery Host** operation.

■ Add Symmetrix, CLARiiON, or VNX enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI, NaviSecCLI, or Navisphere CLI location on this discovery host. This operation returns the list of Symmetrix, CLARiiON, or VNX arrays (local and remote) that are accessible to the host. To configure disaster recovery for applications, select one or more local arrays.

Note: Any physical or virtual machine can be designated as the array control host, including the virtual machine inside VMware ESX server.

■ Add RecoverPoint appliance for replication using the **+ RecoverPoint** operation.

■ Perform add host, add RecoverPoint appliance, and add enclosure operations for the IMS at the disaster recovery data center as well.

Configuring application disaster recovery using EMC RecoverPoint without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC RecoverPoint when the hosts are not a part of a Microsoft failover cluster.

■ Ensure that RecoverPoint consistency groups are setup on the control host for the replication between the primary and remote arrays.

Note: The replicated and primary LUNs must be on different hosts from different data centers.

■ Ensure that EMC RecoverPoint groups are set up for the CRR replication between the primary and Secondary RecoverPoint Appliance.

■ Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read and write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. Windows Storage Space Storage Pool is not supported.

■ Ensure that you have configured application at the production data centre under the Hyper-V Manager and kept the data files on the replicated volumes.

■ Ensure that the respective remote disks are in the offline state on the Hyper-V server at the recovery data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

Once you have performed the necessary configurations, proceed with disaster recovery specific tasks.

Veritas Resiliency Platform configurations:

- Add the host where the RecoverPoint consistency groups are configured, to the Infrastructure Management Server (IMS) using the **+ Add application host** operation.

- Add Symmetrix, CLARiiON, or VNX enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI, NaviSecCLI, or Navisphere CLI location on this discovery host. This operation returns the list of Symmetrix, CLARiiON, or VNX arrays (local and remote) that are accessible to the host. To configure disaster recovery for applications, select one or more local arrays.

Note: Any physical or virtual machine can be designated as the array control host, including the virtual machine inside VMware ESX server.

- Add RecoverPoint appliance for replication using the **+ RecoverPoint** operation.

- Perform add host, add RecoverPoint appliance, and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitations

- If the application hosts are inside the virtual machines, the replicated data disks must be mapped to these hosts in Raw mode only. Virtual disks are not supported.

- Logical grouping of disks, Windows Server Storage space storage pool, is not supported.

Configuring DNS server settings for a data center

Using the Resiliency Platform console, you can configure the DNS settings for the data center.

You can add DNS servers for the data center or remove the settings for servers that were previously added.

To configure DNS server settings for a data center

1 Prerequisites

You must have the following information:

- The IP address of the DNS server

- The name of the domain, and associated credentials.
  Linux Bind: For TSIG authentication, you need the TSIG key and TSIG private files.

Windows DNS: For GSSAPI authentication, you need the user name and keytab file.

- A test host name and IP address for performing a test operation. The test operation validates the specified DNS configuration.

2 Navigate

⚙ **Settings** (menu bar) > **Infrastructure** > **Details View**

Expand the data center > **Access Profile**

Click the **Windows DNS** or **Bind** tab.

DNS servers already added for the data center are listed in the table. You can remove or add a new DNS server.

3 To add a new DNS server for the data center click **+ Add New DNS**.

4 Specify the IP address for the DNS server and select the purpose, either Rehearsal or Production.

5 Add one or more domains for the DNS server:

- Fill in the domain name and the authentication type. For TSIG, browse to the key and private files. For GSSAPI, enter the user name and browse to the keytab file.

- Enter a test host name and IP address and select **Test**. If the test is successful, that is the DNS configuration is validated, the **Add** button is enabled.

- Select **Add**.

6 If you are done adding domains, select **Next**.

7 To remove a DNS server, right-click the required DNS server in the table and select **Remove**.

Configuring resiliency groups for remote recovery

This chapter includes the following topics:

- Understanding the role of resiliency groups in disaster recovery operations
- Managing applications for remote recovery (DR)

Understanding the role of resiliency groups in disaster recovery operations

To perform disaster recovery (DR) operations on virtual machines or applications, they must be configured for disaster recovery as part of a resiliency group, which is the unit of management and control in Veritas Resiliency Platform.

In the configuration wizard for resiliency groups, you apply a service objective to a resiliency group. When you apply the recover hosts service objective, the wizard prompts you for the additional information required for Resiliency Platform to configure the resiliency group for disaster recovery operations.

After disaster recovery configuration on a resiliency group is complete, you can proceed with DR-specific tasks on the resiliency group, such as migrate and take over.

A Virtual Business Service (VBS) lets you further group these resiliency groups in a multi-tier grouping mechanism, and lets you perform controlled start, stop and recovery operations on these resiliency groups.

Managing applications for remote recovery (DR)

To provide disaster recovery protection, you organize applications into a resiliency group and apply the remote recovery service objective. The wizard prompts for the inputs that are needed for the selected service objective and for the replication technology. The wizard then implements the configuration that is required for DR operations.

To manage applications for remote recovery (DR)

1   Prerequisites

Ensure that you have completed the configuration prerequisites before you manage the application for remote recovery.

See "Prerequisites for configuring applications for disaster recovery" on page 43.

2   Navigate

⬛   **Assets** (navigation pane) > **Resiliency Groups** tab > **Manage & Monitor Virtual Machines or Applications**

You can also launch the wizard from the **Unmanaged** or **Overview** tabs.

3   Select the applications:

On the **Select Assets** page, use one or more drop-downs to filter your list of applications.

4   Drag and drop applications to **Selected Instances**.

5   The next page displays the environment for the selected assets.

6   The next page lists the service objectives that are available for the selected applications. You can expand the service objectives to view details. Select the service objective that lets you perform the remote recovery of the applications.

7   Select the target (recovery) data center and then select the target data center application for each source data center application. Click **Next**.

8   On the **Customize Network** page, you can do the following:

■   Choose **Apply production - DNS customization** to customize the DNS setting while performing the migrate or takeover operations.

■   Choose **Apply Rehearsal - DNS customization** to customize the DNS setting while performing the rehearsal operation.

■   Choose to create pointer (PTR) records for the host. A PTR record resolves the IP address to the host name. It is used for reverse DNS lookups.

- Specify whether you want to abort the DR operation if the DNS updates fail.

- For each application specific source IP, select a target IP from the drop down by double clicking on the **Target IP** box.

- Enter DNS registered host name for the application IP by double clicking on the **DNS host name** box.

- Remove the IPs rows for which you do not intend to set DNS customization.

When you complete your selections, click Next

9 Verify the summarized information and enter a name for the resiliency group.

10 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the applications for DR operations. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See "Viewing activities" on page 79.

Verify that the new resiliency group is added to the **Resiliency Groups** tab.

See "Viewing resiliency group details" on page 38.

Rehearsing DR operations to ensure DR readiness

This chapter includes the following topics:

- About ensuring the disaster recovery readiness of your assets
- Rehearse operations for applications - array-based replication
- Performing the rehearsal operation
- Performing cleanup rehearsal

About ensuring the disaster recovery readiness of your assets

Resiliency Platform provides a rehearse operation to help you ensure the disaster recovery readiness of the assets in your protected resiliency groups.

A disaster recovery rehearsal is an operation to verify the ability of your configured resiliency group to fail over on to the target (recovery) data center during disaster. A rehearsal is a zero-downtime test that mimics the configuration, the application data, the storage, and the failover behavior of your resiliency group.

When you are satisfied with the testing of the simulated failover to the target data center, you can use the cleanup rehearsal operation to clean up any temporary objects created during the rehearsal.

Note: If you perform rehearsal or cleanup rehearsal operations on a resiliency group with InfoScale applications, you need to setup the firedrill service group in the GCO.

Rehearse operations for applications - array-based replication

The requirements for rehearse operations for applications depends on the replication type.

Rehearse operations with EMC SRDF based replication:

- Device group should be associated with the snapshot LUNs. Veritas Resiliency Platform supports Timefinder Snap and Timefinder Mirror (BCV).

- Rehearsal operations for resiliency groups that are replicated using EMC SRDF technology in Asynchronous mode cannot be performed using TimeFinder Snap technology (VDEV devices). You need to configure Timefinder Mirrors (BCV devices) to perform the rehearsal operations on such resiliency groups.

- When the rehearse operation is initiated, the Resiliency Platform creates point in time snapshots as a part of the rehearsal operations, since it cannot work with the existing snapshots.

Note: If there are any active snapshots that are in progress, you need to terminate the snapshots and refresh the asset discovery.

- The snapshot disks are enabled on the DR hosts.

- For Linux hosts, the logical volume manager (LVM) volume group (VG) is imported using the snapshot disks. The LVM volume, on the snapshot VG, is then mounted on a mount point. This mount point is same as that of the VG on the DR host if you were to perform the takeover or migrate operation.

- For Windows host, the snapshot disk is assigned a drive letter. This drive letter is the same as that of the replicated disk if you were to perform the takeover or migrate operation.

- The application is then started on the DR hosts. The application starts to consume storage from the snapshot disks instead of the replicated disks.

- During the Rehearse cleanup operation, the above tasks are reversed and the snapshots are terminated, bringing the system back to the original state.

Rehearse operations with NetApp SnapMirror based replication:

- NetApp SnapMirror based replication uses FlexClone for the Rehearse operation and so NetApp storage server must be enabled with the FlexClone license.

- When the rehearse operation is initiated, the Resiliency Platform creates a point in time volume snapshot as a part of the rehearsal operation. The snapshot volume is exported and mounted on the DR host.

> **Note:** Rehearse operation breaks any ongoing replication between the source storage server and the destination storage server as the FlexClone operation cannot be performed on the destination read-only volume. SnapMirror replication resumes after the rehearsal cleanup operation is complete.

- The application is then started on the DR hosts. The application starts to consume storage from the snapshot disks instead of the replicated disks.
- During the Rehearse cleanup operation, the above tasks are reversed and the snapshots are terminated, bringing the system back to the original state.

Performing the rehearsal operation

Use the **Rehearsal** option on the Resiliency Platform console to perform the disaster recovery rehearsal, which verifies the ability of your configured resiliency group to fail over to the disaster recovery (DR) data center during disaster. A rehearsal is a zero-downtime test that mimics the configuration, application data, storage, replication, and the fail over behavior of your resiliency group.

Rehearsal for applications configured for DR with Microsoft Failover Clustering is not supported.

> **Note:** You can perform the Rehearsal operation only on the recovery data center.

To perform the rehearsal operation

1   Prerequisites

- Each type of replication has prerequisites and limitations for the rehearsal operation.
  See "Rehearse operations for applications - array-based replication" on page 60.

2   Navigate

     ⊞   **Assets** (navigation pane)

                     Resiliency Groups

3   Double-click the resiliency group to view the details page. Click **Rehearsal**.

4   Select the target data center and then click **Next**.

Before you perform the rehearsal operation again, you need to ensure that the previous rehearsal is cleaned up by running the Cleanup Rehearsal operation.

Performing cleanup rehearsal

After you have performed the rehearsal operation successfully to verify the ability of your configured resiliency group to fail over on to the disaster recovery data center, you can use the cleanup rehearsal operation to clean up the rehearsal virtual machines or applications in the resiliency group. All temporary objects created during the rehearsal operation are now deleted.

To perform cleanup rehearsal

1   Navigate

　　▣ 　 **Assets** (navigation pane)

　　　　 Resiliency Groups

2   Double-click the resiliency group to view the details page. Click **Cleanup Rehearsal**.

3   Select the target data center, and then click **Next**.

Performing disaster recovery operations

This chapter includes the following topics:

- Migrating a resiliency group of applications

- Taking over a resiliency group of applications

- Performing the resync operation

Migrating a resiliency group of applications

A typical application migration involves the following steps. These steps are performed automatically by the Resiliency Platform as a part of the migrate operation.

- At the primary data center, stop the application and storage.

- Reverse the replication role.

- At the recovery data center, start the storage and application.

- Update the DNS resource records.

To migrate applications

1 Navigate

Assets (navigation pane)

Resiliency Groups

2 Double-click the resiliency group to view the details page. Click **Migrate**.

3 Select the target data center and click **Next**.

Taking over a resiliency group of applications

Takeover is an activity initiated when the production data center is down due to any disaster or natural calamities, and the applications need to be restored at the recovery data center to provide business continuity.

To perform takeover operation on applications

1 Navigate

	Assets (navigation pane)

   Resiliency Groups

2 Double-click the resiliency group to view the details page. Click **Takeover**.

3 Select the target data center and click **Next**.

Performing the resync operation

When disaster strikes on a production data center, the takeover operation is invoked to start the resiliency groups on the recovery data center.

Since the production data center is not working, the data replication between the two sites does not happen. After the production site is back up and running, you need to prepare the production site for the next failover or for a migration operation. This preparation includes cleaning up any residue and resuming the replication from the recovery to the production site.

Use the Resync operation on the Resiliency Platform console to automate these steps for the required resiliency groups. This operation cleans up the residue which includes stopping applications and virtual machines, unregistering virtual machines, unmounting file systems, datastores, etc.

Performing the resync operation

1 Navigate

Assets (navigation pane)

Resiliency Groups

2 Double-click the resiliency group to view the details page. Click **Resync**.

3 In the **Resync** panel, select the production data center name from the drop-down list, and click **Next**.

Monitoring and reporting on assets status

This chapter includes the following topics:

- About the Resiliency Platform Dashboard
- Understanding asset types
- Displaying an overview of your assets
- Viewing reports

About the Resiliency Platform Dashboard

The Resiliency Platform Dashboard gives you an overview of your resiliency domain. Use the Dashboard to answer questions such as the following:

- Which of my data centers have Resiliency Platform managed assets?
- What is the mix of my assets by type and platform?
- Which assets are configured for disaster recovery?

The Dashboard has the following areas:

Global View	A world map that identifies the data centers that contain Resiliency Platform managed assets.
	Lines between data centers indicate that replication takes place between the locations.
	Mouse over an icon for basic Resiliency Platform configuration and asset configuration information for that data center. Click **More** for detailed information and recent activity.
Resiliency Groups and **Virtual Business Services** summaries	The upper right section of the dashboard displays total number of resiliency groups and virtual business services in the resiliency domain, as well as those at risk and normal.
	Click a square in either the **Resiliency Groups** or **Virtual Business Services** summary to display a tab of detailed information.
Virtual Machines by Platform and OS	Displays a summary of virtual machines in all data centers or information on a single data center. Use the drop-down list to filter your results. The summary lists the virtual machine types by percentage and the platform types by number.
Application environment	Displays the number of applications and the application types. The chart shows the number of applications that are managed by InfoScale and those that are not managed by InfoScale.
Applications by Type	Displays a summary of application types in all data centers or in a single data center. Use the drop-down list to filter your results.
Top Resiliency Groups by Replication Lag	Ranks the resiliency groups according to how long it takes the recovery data center to be in sync with the active data center.
By Service Objective	Displays the percentage of virtual machines and applications that are unprotected or unmanaged.
	Use the drop-down list to filter your results.

See

Understanding asset types

On the Resiliency Platform console Assets page, assets are classified as follows.

Asset	Description
Resiliency Group	A group of applications or virtual machines under Resiliency Platform control. You can use Resiliency Platform to start and stop the resiliency group, as well as protect and manage it.
	The Overview tab identifies whether or not resiliency groups are protected. An unprotected resiliency group is one that is configured to support monitoring and start and stop operations only. A protected resiliency group supports data recovery operations as well.
Virtual Business Service	A collection of resiliency groups logically grouped for a specific business purpose.
Unmanaged	An application or virtual machine that Resiliency Platform discovers in your environment, but that is not under Resiliency Platform management. You cannot use any Resiliency Platform features with these assets until they become a part of a resiliency group.

Displaying an overview of your assets

The **Assets** page gives you an overview of all your resiliency groups and virtual business services (VBSs). You can also click links on the page to create resiliency groups and VBSs.

To access the **Assets** page, go to the navigation pane on the left side of the screen, and click:

The **Assets** page is organized into the following categories:

- Unprotected resiliency groups, are groups under Resiliency Platform control, but that do not have disaster recovery configured.
  See "Managing applications for basic monitoring" on page 31.

For unprotected and protected resiliency groups, the screen also displays the following:

- The number of resiliency groups that are based on virtual machines and the number that are based on applications

- The number of unmanaged virtual machines or applications; that is, the assets that Resiliency Platform is aware of but that are not managed or protected in resiliency groups.

For VBSs, the screen displays the following:

- The number of VBSs that are created from virtual machines and the number that are created from physical assets.

- The number of resiliency groups within the VBSs that are protected and the number that are only managed (not protected).

Viewing reports

Veritas Resiliency Platform provides a console for viewing the following reports:

Resiliency Groups and VBS Summary	Provides details about the resiliency groups and VBSs in the data centers across all sites.
Activity Distribution History	Provides information about tasks, such as migrate, takeover, rehearse, start, and stop, performed for a specified duration.
Recovery Activity History by RG	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, for each resiliency group.
Recovery Activity History by VBS	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, for each VBS.

To view a report

1 Navigation

Click **Reports** (menu bar).

2 Do one of the following:

- Click **Run** to receive the report on the specified email address in HTML or PDF format, or as a comma separated (.CSV) file. You can also view the saved report on the console.

- Click **Schedule** to create a report generation schedule.

For more information on configuring email settings and scheduling reports, refer to the *Deployment Guide*.

Monitoring risks

This chapter includes the following topics:

- About risk insight

- Displaying risk information

- Predefined risks in Resiliency Platform

- Viewing the current risk report

- Viewing the historical risk report

About risk insight

The objective of the Risk Insight feature is to notify you about the vulnerabilities that might impact the recoverability or continuity of your protected assets.

Risk Insight detects the changes to the state and configuration of your protected assets. It identifies if there is a risk to the recoverability or continuity of your protected assets.

Veritas Resiliency Platform also enables you to set up the replication lag threshold or service level threshold. Risk insight alerts you when the replication lags beyond the threshold that you specified.

Risk insight generates two types of reports:

- **Current risk reports**: Provides the summary and detail information about all the current risks in your data center.

- **Historical risk reports**: Provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

These reports help you take actions to prevent such risks. The historical risk data is purged after a period of two years.

The risks covered by risk insight can be classified into three main categories:

Table 10-1

Risk Type	Description
Recoverability	Risks that may impact the ability to recover and run the application on the recovery site.
Continuity	Risks that may impact the ability to run your applications without disruption either on your production site or on your recovery site.
SLA	Risks that may impact the ability to fulfill the service level agreements (SLA) for your applications.

On the basis of criticality, the risks can be classified into two types:

Table 10-2

Risk type	Description
Error	A risk that disrupts any stated goals of the product. An error must be fixed to make the product work as expected.
Warning	A risk that jeopardizes any stated goals of the product. A warning alerts you about a potential problem in your environment.

Displaying risk information

Resiliency Platform identifies and flags several risks that may occur during data center operations. Some of these risks are transient. They are temporary and resolve themselves without your intervention. Other risks require intervention and troubleshooting to resolve.

You can display risks in the following ways:

Table 10-3   Ways to display risks

To display ...	Do the following:
A complete list of risks across the resiliency domain	**1** On the menu bar, select ⊞ **More Views** > **Risks** **2** On the **Risk** page, double-click a risk in the table to display detailed information.

Table 10-3   Ways to display risks *(continued)*

To display ...	Do the following:	
Risks that are associated with a specific resiliency group or virtual business service	**1**	On the navigation pane, select (Assets) and the tab for either **Resiliency Groups** or **Virtual Business Services**.
	2	On the tab, double-click a resiliency group or virtual business service to display detailed information.
	3	On the details page, note any risks that are listed in the **At Risk** area, and double-click the risk for details.

In addition to the above mentioned views, the **More views** > **Logs** > **All** view and the **More views** > **Logs** > **Notification** view also includes the notification about the risks in your environment. You can double-click any row to view the detailed description of the error and suggested resolution for the error.

Predefined risks in Resiliency Platform

Table 10-4 lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 10-4   Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Veritas Infoscale Operations Manager disconnected	Checks for Veritas Infoscale Operations Manager to Resiliency Manager connection state	1 minute	Error	All operations	Check Veritas Infoscale Operations Manager reachability Try to reconnect Veritas Infoscale Operations Manager

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Password Incorrect	Checks if vCenter password is incorrect	5 minutes	Error	- On primary site: start or stop operations - On secondary site: migrate or takeover operations	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown.	Real time, when resiliency group is created	Error	- Migrate - Stop	- In case of VMWare, install VMWare Tools - In case of Hyper-V, install Hyper-V Integration Tools
Snapshot removed from Virtual Machine	Checks if snapshot has been removed from virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Remove and re-add the virtual machine to the Resiliency group by editing Resiliency group
Data Mover Daemon Crash	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Snapshot created on Virtual Machine	Checks if a snapshot has been created on Virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
DataMover virtual machine in noop mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Resiliency group configuration drift	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	■ Migrate ■ Resync	Edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment.	Real time	Warning	■ Migrate ■ Takeover	Edit the resiliency group or add a Global user
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager.	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state.	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue

Table 10-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
DNS removed	Checks if DNS is removed from the resiliency group where DNS customization is enabled	real time	Warning	▪ Migrate ▪ Takeover	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the VM has not been restarted after add host operation	2 hours	Error	Configure DR	Restart the VM after add host operation
New VM added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group.	5 minutes	Error	▪ Migrate ▪ Takeover ▪ Rehearsal	Add the virtual machine to the resiliency group.
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center.	5 minutes	Warning	▪ Migrate ▪ Takeover	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group.	5 minutes	Error	▪ Migrate ▪ Takeover	Contact the enclosure vendor.

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Remote mount point already mounted	Checks if the mount point is not available for mounting on target site for any of the following reasons: ■ Mount point is already mounted. ■ Mount point is being used by other assets.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover	Unmount the mount point that is already mounted or is being used by other assets.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover ■ Rehearsal	Delete or move some files or uninstall some non-critical applications to free up some disk space.
ESX not reachable	Checks if the ESX server is in a disconnected state.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the ESX server connection issue.
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the virtualization server connection issue. In case of a password change, resolve the password issue.

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment.	6 hours	Warning	■ Migrate ■ Takeover	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target.
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center.	30 minutes	Error	Migrate	Check the following and fix: ■ Host is up on recovery data center. ■ Host is accessible from recovery datacenter IMS. ■ Time is synchronized between host and recovery datacenter IMS.

Table 10-4   Predefined risks *(continued)*

Viewing the current risk report

This report provides the summary and detail information about all the current risks in your data center. The high-level summary shows the total number of risks and its distribution by severity.

The **Distribution by type** chart displays the severity-wise distribution for recoverability, continuity, and service level agreement (SLA).

The **Unresolved risks** chart shows the risks that are unresolved for more than one month, for last one month, and for last one week. The **Recent Risks** chart shows the recent risks that are generated in the last 24 hours.

The **Current risks details by type** table provides detailed information such as the name of the resiliency group which is at risk, the name of the risk, its description, object at which the risk is generated, severity, and date and time on which the risk was generated.

To view the current risk report

1 Navigation:

Click **Reports** (menu bar).

2 In the **Risk** > **Current Risk Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Viewing the historical risk report

This report provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

The high-level summary shows the total number of risks and its distribution by the time the risks have been open. The information is categorized under various headings such as **Carried forward**, **New**, **Closed**, and **Still open**.

Within these categories, you can see severity wise distribution (high or low) and category wise distribution (recoverability, continuity, and service level agreement) of the risks.

The detailed analysis is displayed in the form of various charts:

■ The various charts under **Risk by Category** display the open risks and new risks in the recoverability, continuity, and SLA categories at specific points of time within the duration specified by you.

■ The **Resolving time chart** shows the average time to resolve the risk within the recoverability, continuity, and SLA categories.

■ The **5 risks that took the longest time to resolve** chart shows the top 5 risks that took the longest time to be resolved, within the recoverability, continuity, and SLA categories. This information is displayed per resiliency group or per Virtual Business Service (VBS).

To view the historical risk report

1 Navigation:

Click **Reports** (menu bar).

2 In the **Risk** > **Risk History Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Managing activities and resiliency plans

This chapter includes the following topics:

- Managing activities
- Managing resiliency plans

Managing activities

Using the Veritas Resiliency Platform console, you can view the sub task information for a task or an operation that is performed on the console.

See "Viewing activities" on page 79.

See "Aborting a running activity" on page 80.

Viewing activities

Using the Veritas Resiliency Platform console, you can view the sub task information for a task or an operation that is performed on the console. You can view the details on the **Activities** page. Details such as the status of the operation (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed are displayed. You can view these details for a currently running task and for the completed tasks. On the **Current** page you can abort a running task.

Click on a currently running task, to view the details in a graphical representation. The steps that are completed are shown in green color along with the success icon. The ongoing steps are in blue color with the loader image, and the future steps are in gray. Expand **Execution Details** to view all the sub-tasks that comprise the task.

To view activities

1 Navigate

⊞ **Activities** (menu bar).

2 Choose either of the following:

- Select **Current** to view the currently running tasks.
- Select **Completed** to view the historical tasks.

To view recent activities, click **Recent Activities** on the bottom pane.

See "Aborting a running activity" on page 80.

Aborting a running activity

Using the Veritas Resiliency Platform console, you can abort a task or an operation which is currently running. You can abort an operation that is executed using a resiliency plan or from the console. When you abort an operation, the sub task which is in progress is completed and then the process is aborted. The status of the sub tasks which were already completed does not change.

For example, the migrate resiliency group operation has six sub tasks. If you abort the operation while the first sub task, Stop Virtual Machine, is in progress, then the Stop Virtual Machine sub task is completed and the remaining sub tasks are skipped. If you restart the migrate operation, it starts from the beginning.

To abort an activity

1 Navigate

Do one of the following:

⊞ **Activities**. Skip to 2

Recent Activities (bottom pane). Click **Abort** on the required activity.

2 In the **Current** activities page, place your cursor on the activity that you want to abort. Do one of the following:

- Right click your mouse and click **Abort**.
- Click on the vertical ellipsis and select **Abort**

See "Viewing activities" on page 79.

Managing resiliency plans

Veritas Resiliency Platform provides a console for creating and customizing resiliency plans. The following topics cover how to create, edit, delete resiliency plan templates and resiliency plans and how to execute resiliency plans.

About resiliency plans

Using the Veritas Resiliency Platform console you can create customized resiliency plans. A resiliency plan is a customized set of tasks that you can run as a single operation. You add each task and the particular assets on which to run the task. If you intend to use the same sequence of tasks on different assets, you can create a resiliency template. You can save the template and use it to create multiple resiliency plans.

For example, you can create a resiliency plan template to migrate a resiliency group. Then you can add a resiliency group to the template to create a plan. You can create multiple plans using the same template.

You can create customized resiliency plans for performing all the disaster recovery operations such as migrate, takeover, rehearsal, cleanup rehearsal, and resync. You can also create customized resiliency plans for executing a manual task or a custom script.

You do not have to create a template in order to create a resiliency plan. Resiliency plans can be created using blank templates.

Note: To create a plan for migrate, takeover, rehearsal, or cleanup rehearsal operation, configure disaster recovery task must be successful on the selected resiliency group.

You can schedule the resiliency plan to run at a particular time.

Using these predefined templates, you can create resiliency plans by adding assets to the template. You can then run these plans on a later date.

In case of InfoScale environment in Veritas Resiliency Platform 2.0 update 2.0.0.100, you need to enable script execution on the hosts. If you create a resiliency plan in Veritas Resiliency Platform 2.0 and then apply update 2.0.0.100, you are prompted to enable the script execution on the hosts when you edit or run the plan.

See "Creating a new resiliency plan template" on page 82.

See "Creating a new resiliency plan" on page 87.

Creating a new resiliency plan template

Using the Veritas Resiliency Platform console, you can create a customized resiliency plan template for the following operations:

- Start and stop a resiliency group.

- Migrate and takeover a resiliency group.

- Manual task
  See "About manual task" on page 83.

- Run a custom script
  See "About custom script" on page 84.

To create a template, you need to drag and drop the required operation from the stencil into the canvas below. The arrow lets you connect various operations in the canvas.

For example, if you want to create a template to perform the Start Resiliency Group task, drag the operation from the top bar into the canvas. Now click on the arrow on the **Start** action box and drag the mouse to the **Start Resiliency** In addition to the above listed tasks, you can also add a custom script Manual task in the resiliency plan. This task temporarily pauses the operation letting you perform a task before proceeding further.

Group action box. Similarly you can drag the arrow from the **Start Resiliency Group** action box to the **End** action.

To create a new resiliency plan template

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Templates** section, click **New**.

3   In the **Create New Template** wizard panel, enter a name and a description for the template.

4   Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.

5   Click **Create**.

See "About resiliency plans" on page 81.

About manual task

Using the Resiliency Platform console, you can add a manual task in the resiliency plan. The purpose of including this task in resiliency plan is to temporarily pause the operation of the resiliency plan to perform a task or validate a step before you proceed further.

You can specify a timeout for the manual task. After the specified timeout expires, the manual task in the resiliency plan is marked as complete and the resiliency plan proceeds further.

Alternatively, you can opt for manually resuming the process. In this case, the resiliency plan enters into a pause state. You need to go to the **Inbox** in Resiliency Platform console and click **Resume** on the corresponding entry in the **Inbox**. You can also resume the resiliency plan by right-clicking the corresponding entry in **Activities > Current Activities** and selecting **Resume**.

Using manual tasks in resiliency plans

Using the Resiliency Platform console, you can add a manual task in the resiliency plan.

To use a manual task in a resiliency plan

1   You can add a manual task to a resiliency plan template or to a resiliency plan.

See "Creating a new resiliency plan template" on page 82.

See "Creating a new resiliency plan" on page 87.

2   Drag and drop **Manual Task** into the canvas. Click the pencil icon in the action box to add the task details.

3   Provide a name for the manual task.

4   Describe the reason why you want to add this manual task to the resilient plan.

5   Select your choice for resuming the process manually or automatically. If you select the option for automatically resuming the process after a timeout, enter the duration of timeout in minutes. Click **Save**.

About custom script

Using the Resiliency Platform console, you can add a custom script execution task in the resiliency plan. You can use the custom script execution task to perform customized operations before executing the next step of the resiliency plan such as repurposing capacity on the recovery site, orchestrate network changes, or any kind of post-processing.

Custom Script execution requires Resiliency Platform deployed on the Resiliency Manager, Infrastructure Management Server (IMS) and the hosts executing custom scripts. In addition, if you are using Resiliency Platform with Veritas InfoScale, the Veritas Resiliency Platform Enablement add-on has to be manually installed on applicable hosts.

The custom script can be in any format that can be directly executed on a shell on the target host. For the Linux hosts, it may be an executable or a script that specifies the interpreter on the shebang line such as a shell or a Perl script. For Windows hosts, it may be an executable or a script with known extension such as a bat file or a PowerShell script. The Script is executed as root user on a UNIX host or as Local System on a Windows host. You may use `sudo` or `RunAs` commands to execute some other scripts from these custom scripts.

Before you can execute the script as part of the resiliency plan, you need to manually copy the script to the *VRTSsfmh InstallDir*/vrp/scripts directory on the host.

Where, *VRTSsfmh InstallDir* is /opt/VRTSsfmh on the Unix/Linux hosts and *SystemDrive*/Program Files/VERITAS/VRTSsfmh on the Windows hosts. Copying the script to these specific folders enforces the security policy for running a custom script since these folders can be accessed only by a root user or a Local System.

Exit code from script execution determines the success or failure of the task in the resiliency plan workflow. An exit code of zero means the script execution was successful while a non-zero exit code means the script execution failed. If you select the option to ignore the exit code, the script task is always marked as successful after completion of the script. You can select this option, if your script does not return any exit code. You can view the output of the script in activity details for the resiliency plan in Resiliency Platform console.

If you uninstall the host package from the host where you have copied your custom script, the custom script is removed from the host as part of the uninstallation process.

Using custom scripts in resiliency plans

Using the Resiliency Platform console, you can add a custom script execution task in the resiliency plan.

To use a custom script execution task in a resiliency plan

1   You can add a custom script execution task to a resiliency plan template or to a resiliency plan.

    See "Creating a new resiliency plan template" on page 82.

    See "Creating a new resiliency plan" on page 87.

2   Drag and drop **Custom Script** into the canvas. Click the pencil icon in the action box to add the task details.

3   Enter a name for the custom script.

4   Select the data center and the host where you want to execute the script. Click **Next**.

5   Enter the following details:

    ■   The relative path of the script on the specified host. The script path that you enter is taken as relative to the *VRTSsfmh InstallDir*/vrp/scripts/ directory path.
        For example, if you enter the path of the script as
        myscripts/backup_scripts/*script_name*, then the complete path considered by the system will be *VRTSsfmh InstallDir*/vrp/scripts/myscripts/backup_scripts/*script_name*.

    ■   Command-line arguments to the script. This is an optional input field.

    ■   Timeout for the script. By default, there is no timeout for the script execution. You can specify a timeout for the script execution. After the specified timeout expires, the script execution task in the resiliency plan is marked as failure but the script execution task is not stopped. The script execution may continue in the background. If you do not specify any timeout, the task will wait till the script is not completed.

6   Click **Save**.

Editing a resiliency plan template

Using the Veritas Resiliency Platform console, you can edit an existing resiliency plan template.

You can add assets to these templates and create a customized resiliency plan. Any changes to the template do not affect the existing resiliency plans that you created from the template.

To edit a resiliency plan template

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Templates** list, place your cursor on the row which you want to edit. Do one of the following:

- Right click your mouse and click **Edit**.

- Click on the vertical ellipsis and select **Edit**.

3 In the **Edit Template** wizard panel, edit the required actions and click **Save**.

The steps for editing the plan are the same as creating it.

See "Creating a new resiliency plan template" on page 82.

Deleting a resiliency plan template

Using the Veritas Resiliency Platform console you can delete an existing resiliency plan template.

Deleting the template does not affect the existing resiliency plans that you created from the template.

To delete a resiliency plan template

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Templates** list, place your cursor on the row which you want to delete. Do one of the following:

- Right click your mouse and click **Delete**.

- Click on the vertical ellipsis and select **Delete**.

3 In the **Delete Template** panel click **Delete**.

See "Creating a new resiliency plan template" on page 82.

Viewing a resiliency plan template

Using the Veritas Resiliency Platform console, you can view the details of a resiliency plan template. To view the details of the resiliency plan templates, you need to have at least guest persona assigned to you.

To view a resiliency plan template

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Templates** list, do one of the following:

- Double click the row that you want to view.

- Select the row that you want to view, right click and select Details.

- Select the row that you want to view, click on the vertical ellipsis and select Details.

3   You can now view the details of the resiliency plan template.

Creating a new resiliency plan

Using the Veritas Resiliency Platform console, you can create a new resiliency plan for the following operations. Resiliency plans can be created using an existing template or with a blank template. When you create a plan using a blank template, you need to create the plan and add the assets at the same time.

- Start and stop a resiliency group.

- Rehearsal and cleanup rehearsal of a resiliency group.

- Migrate and takeover a resiliency group.

- Manual task
  See "About manual task" on page 83.

- Run a custom script
  See "About custom script" on page 84.

Note: To create a plan for migrate, takeover, rehearsal, or cleanup rehearsal operation, disaster recovery must be configured successfully on the selected resiliency group or the VBS.

To create a new resiliency plan using blank template

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** section, click **New**.

3   In the **Create Saved Plan - Select Template** wizard panel, select **Blank Template**, and click **Next**.

4   In the **Add Assets** panel, enter name and description.

5 Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.

6 Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.

7 Click **Submit**.

To create a new resiliency plan using predefined template

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Saved Plans** section, click **New**.

3 In the **Create Saved Plan - "Select Template"** wizard panel, select **Pre-defined Template**.

4 Select a template from the list and click **Next**.

5 In the **Add Assets** panel, name and description are pre-populated.

6 Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.

7 Click **Submit**.

See "About resiliency plans" on page 81.

See "Deleting a resiliency plan" on page 89.

See "Executing a resiliency plan" on page 89.

Editing a resiliency plan

Using the Veritas Resiliency Platform console, you can edit a resiliency plan.

To edit a resiliency plan

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Saved Plans** list, place your cursor on the row which you want to edit. Do one of the following:

- Right click your mouse and click **Edit**.

- Click on the vertical ellipsis and select **Edit**.

3 In the **Edit Saved Plan** wizard panel, edit the required actions and click **Submit**.

The steps for editing the plan are the same as creating it.

See "Creating a new resiliency plan" on page 87.

Deleting a resiliency plan

Using the Veritas Resiliency Platform console, you can delete a resiliency plan.

To delete a resiliency plan

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Saved Plans** list, place your cursor on the row which you want to delete. Do one of the following:

- Right click your mouse and click **Delete**.

- Click on the vertical ellipsis and select **Delete**.

3 In the **Delete Saved Plan** panel click **Delete**.

See "Creating a new resiliency plan" on page 87.

Executing a resiliency plan

Using the Veritas Resiliency Platform console, you can execute a resiliency plan. After executing the resiliency plan, you can navigate to the **Activities** page to view the progress of the plan.

To execute a resiliency plan

1 Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2 In the **Saved Plans** list, place your cursor on the row which you want to execute. Do one of the following:

- Right click your mouse and click **Execute**.

- Click on the vertical ellipsis and select **Execute**.

3 In the **Execute Saved Plan** panel click **Execute**.

See "Creating a new resiliency plan" on page 87.

Viewing a resiliency plan

Using the Veritas Resiliency Platform console, you can view the details of a resiliency plan. To view the details of the resiliency plans, you need to have at least guest persona assigned to you.

You can also launch operations such as edit a resiliency plan or delete a resiliency plan from this view.

See

See

To view a resiliency plan

1   Navigate

   Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

   ■   Double click the row that you want to view.

   ■   Select the row that you want to view, right click and select **Details**.

   ■   Select the row that you want to view, click on the vertical ellipsis and select **Details**.

3   You can now view the details of the resiliency plan. Click the watch icon to see the details of the components of a resiliency plan such as a custom script or a manual task.

Creating a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can create a schedule for a resiliency plan.

To create a schedule for a resiliency plan

1   Navigate

   Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

   ■   Double click the row for which you want to create a schedule. In the **Schedule** section of details page, click **New**.

   ■   Select the row for which you want to create a schedule, right click and select **Create Schedule**.

   ■   Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Create Schedule**.

Editing a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can edit a schedule for a resiliency plan.

To edit a schedule for a resiliency plan

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to edit a schedule. In the **Schedule** section of details page, click **Edit**.

- Select the row for which you want to create a schedule, right click and select **Edit Schedule**.

- Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Edit Schedule**.

Deleting a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can delete a schedule for a resiliency plan.

To delete a schedule for a resiliency plan

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to delete a schedule. In the **Schedule** section of details page, click **Delete**.

- Select the row for which you want to edit a schedule, right click and select **Delete Schedule**.

- Select the row for which you want to edit a schedule, click on the vertical ellipsis and select **Delete Schedule**.

Viewing a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can view a schedule for a resiliency plan. To view the details of the resiliency plans, you need to have at least guest persona assigned to you.

You can also launch operations such as edit a schedule or delete a schedule from this view.

See "Editing a schedule for a resiliency plan" on page 90.

See "Deleting a schedule for a resiliency plan" on page 91.

To view a schedule for a resiliency plan

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to view a schedule.

- Select the row for which you want to view a schedule, right click and select **Details**.

- Select the row for which you want to view a schedule, click on the vertical ellipsis and select **Details**.

3   In the **Schedule** section of details page, view the details of the schedule.

Troubleshooting

This appendix includes the following topics:

- Viewing events and logs in the console
- Displaying risk information

Viewing events and logs in the console

Veritas Resiliency Platform maintains the following types of logs that can be viewed in the web console:

System logs: System logs are typically the result of a user performing an operation in the console.

Audit logs: Audit logs are primarily used for security audits. They leave a chronological trail of activities performed on the system. They identify user, activity, affected objects, etc. They help track the individuals responsible for activities and detect security violations.

Event and notification logs: Event and notification logs are not necessarily related to user activity; they can include information such as a server going down. Events can be public or private. Rules can be configured to notify users by email of selected public events. Private events are typically unrelated to user-initiated operations. Private events are displayed in the console for troubleshooting but are not available to include in rules for notification.

By default, logs and SNMP traps are retained for 2 years. This retention period can be modified in the product settings in the console.

To view events and logs

1 Navigate

 More Views (menu bar) > **Logs**

 You can also view new notifications from the **Notifications** icon.

2 To view logs by type (System, Audit, or Notification) select the appropriate tab. You can filter by the product service and by severity (information, warning, or errors) or type (public, private), depending on the tab.

Displaying risk information

Resiliency Platform identifies and flags several risks that may occur during data center operations. Some of these risks are transient. They are temporary and resolve themselves without your intervention. Other risks require intervention and troubleshooting to resolve.

You can display risks in the following ways:

Table A-1   Ways to display risks

To display ...	Do the following:
A complete list of risks across the resiliency domain	**1** On the menu bar, select **More Views** > **Risks** **2** On the **Risk** page, double-click a risk in the table to display detailed information.
Risks that are associated with a specific resiliency group or virtual business service	**1** On the navigation pane, select (Assets) and the tab for either **Resiliency Groups** or **Virtual Business Services**. **2** On the tab, double-click a resiliency group or virtual business service to display detailed information. **3** On the details page, note any risks that are listed in the **At Risk** area, and double-click the risk for details.

In addition to the above mentioned views, the **More views** > **Logs** > **All** view and the **More views** > **Logs** > **Notification** view also includes the notification about the risks in your environment. You can double-click any row to view the detailed description of the error and suggested resolution for the error.

activity	A task or an operation performed on a resiliency group.
add-on	An additional software package that can be installed on hosts by the Infrastructure Management Server (IMS) for specialized uses.
asset infrastructure	The data center assets that can be added to the Infrastructure Management Server (IMS) for IMS discovery and monitoring. For example, virtual machines or virtualization servers.
assets	In Veritas Resiliency Platform, the virtual machines or applications that have been discovered by the Infrastructure Management Server (IMS) and that can be grouped into resiliency groups.
CLISH	Command Line Interface SHell. Provides the command line menu on the virtual appliance for use after the initial bootstrap configuration.
data center	A location that contains asset infrastructure to be managed by Veritas Resiliency Platform.
	For the disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.
host	Physical servers, virtual machines, or Hyper-V servers that are added to the Infrastructure Management Server (IMS) as hosts.
	Adding the assets as hosts installs the host package that is used by the IMS for discovery and monitoring.
Infrastructure Management Server (IMS)	The Veritas Resiliency Platform component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager.
migrate	A planned activity involving graceful shutdown of virtual machines at the production data center and starting them at the recovery data center. In this process, replication ensures that consistent virtual machine data is made available at the recovery data center.
persona	A user role that has access to a predefined set of jobs (operations). Used to assign permissions to users and groups for Veritas Resiliency Platform web console operations.
product role	The function configured for a Veritas Resiliency Platform virtual appliance.

	For example, a virtual appliance can be configured as a Resiliency Manager, Infrastructure Management Server (IMS) or both.
production data center	The data center that is normally used for business. See also recovery data center.
recovery data center	The data center that is used if a disaster scenario occurs. See also production data center.
rehearsal	A zero-downtime test that mimics the configuration, application data, storage, and the failover behavior of the resiliency group.
	Rehearsal verifies the ability of the resiliency group to fail over to the recovery data center during a disaster.
resiliency domain	The logical scope of a Resiliency Platform deployment. It can extend across multiple data centers.
resiliency group	The unit of management and control in Veritas Resiliency Platform. Related assets are organized into a resiliency group and managed and monitored as a single entity.
Resiliency Manager	The Veritas Resiliency Platform component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console.
resiliency plan	A collection of tasks or operations, along with the relevant assets, which are performed in a predefined sequence.
resiliency plan template	A template defining the execution sequence of a collection of tasks or operations.
take over	An activity initiated by a user when the production data center is down due to a disaster and the virtual machines need to be restored at the recovery data center to provide business continuity.
tier	Within a virtual business service (VBS), resiliency groups are arranged as tiers. Tiers represent the logical dependencies between the resiliency groups and determine the relative order in which the resiliency groups start and stop.
virtual appliance	An appliance that includes the operating system environment and the software application which are deployed together as a virtual machine.
	The Veritas Resiliency Platform virtual appliance is deployed as a virtual machine and then configured with basic settings and a role (for example, Resiliency Manager).
virtual business service (VBS)	A multi-tier IT service where each VBS tier hosts one or more resiliency groups. A VBS groups multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. You can also migrate/takeover the entire VBS.
web console	The web-based management console on the Resiliency Manager that is used to configure the settings for the resiliency domain and perform operations.

Index