Dell™ vRanger™ 7.2

User's Guide

Legend

⚠ **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

⚠ **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

ⓘ **IMPORTANT NOTE**, **NOTE**, **TIP**, **MOBILE**, or **VIDEO:** An information icon indicates supporting information.

Contents

Introduction to this guide

vRanger – at a glance

Dell™ vRanger software provides high-speed backup and recovery of VMware®, Microsoft® Hyper-V®, and physical Windows Server® environments. It also provides high-speed replication for VMware. It protects entire virtual environments in minutes, detecting and backing up new VMs automatically, and delivers safe, scalable data protection to even the largest VMware and Hyper-V environments. With vRanger, you can locate and restore individual files in seconds — even if they are buried in virtual and physical backups — from a single, intuitive interface.

vRanger capabilities include:

- Protect entire VMware and Hyper-V environments in minutes
- Deliver high-speed, storage-thrifty backup/restore for VMware, Hyper-V and physical Windows servers
- Quickly locate and restore individual files from virtual and physical backups
- Replicate key VMs for offsite disaster recovery preparedness
- Deploy scalable data protection for even the largest virtual infrastructures

Key benefits

- Exploits vSphere 6 performance and scalability enhancements
- Provides high-speed, resource-efficient backup, replication and recovery of VMware and Hyper-V virtual machine (VM) images
- Provides high-speed backup and recovery of physical servers.
- Delivers up to 90 percent backup storage savings when paired with Dell™ DR Series appliance.
- Includes capability for distributed segment processing through DD Boost™ support for faster backups to EMC® Data Domain® systems
- Optionally deploys as a low-resource consumption vRanger Virtual Appliance (VA) for low-impact scalability
- Offers VA cluster support and easy VA deployment through a wizard
- Improves VM backup and recovery performance with HotAdd support
- Optimizes use of critical resources in virtual environments, including hypervisor hosts, networks and storage

- Protects real-world VM deployments with multiple hosts and hundreds of VMs

- Ensures archive integrity

- Simplifies portability with image backup files

- Delivers proven reliability

- Provides unmatched performance and scalability

- Comes with world-class service and support

Major feature list

- **VMware ReadyTM Certified for vSphere 6** — Ensures reliable, safe, and scalable operation with vSphere 6.

- **VMware ESXi Support** — Provides a wide range of backup and recovery and disaster recovery (DR) capabilities for VMware virtual infrastructures.

- **Microsoft Hyper-V Support** — Protects and provides disaster recovery capabilities for Hyper-V virtual machines.

- **Physical Server Support** - Supports backup and recovery of Windows physical servers, files and folders

- **Virtual Appliance (VA) Architecture** — Through centralized and wizard-driven deployment and administration from the vRanger server, delivers highly scalable and cost-effective distributed data handling and throughput through VAs.

- **Disk-to-Disk Backup and Deduplication** — Supports Dell DR RDA for optimized deduplication and replication

- **EMC Data Domain Boost Support** — Supports EMC Data Domain Boost for optimized deduplication and replication.

- **VMware HotAdd Support** — Performs LAN-free backups with vRanger installed inside a VM and from the vRanger VA. Additionally, HotAdd accelerates network backups of ESXi.

- **Patented Active Block Mapping (ABM)**— Eliminates inactive and white space blocks from protected Windows VMs to speed backup, replication, and recovery jobs as well as reduce network and storage requirements.

- **Instant File-Level Recovery (FLR) for Windows and Linux** — Lets you quickly restore a single file from a backup image in the repository through a one-step process.

- **Native, Full Catalog Capability (Patent-Pending)** — Provides a native, full catalog of every image in the backup repository, enabling immediate identification of available recovery positions, with one-click restore. Wildcard scanning feature quickly locates backup repository files to be restored.

- **Advanced Encryption Standard (AES)-256** — Secures protected images block-by-block on the VMware host as they're read so they're also secure over the network and in the backup repository.

- **Synthetic Recovery** — Delivers single-pass restore, reading each required block only one time from multiple full, incremental, and differential backup images in the repository for the fastest, most efficient results.

- **vMotion and Storage vMotion Support** — vSphere vMotion support ensures vRanger automatically protects VMware VMs as they move from one host to another, even when backup jobs are running. It also ensures vRanger follows VM storage disks when they are relocated to different data stores; locks VM storage disks when vRanger accesses the disks during a job.

- **Advanced Savepoint Management** — Lets you manage and use multiple point-in-time copies of backup and replica images for precise image, file, and object restores.

- **Dynamic Resource Management** — Optimizes real-time use of critical resources; ensures efficiency and that jobs complete faster without exceeding resource capacity.

Target audience

This guide is intended for backup administrators and other technical personnel who are responsible for designing and implementing a backup strategy for the organization. A good understanding of the operating system (OS) on which vRanger is running is assumed.

Recommended additional reading

The following documentation is also available:

- *Dell™ vRanger™ Installation/Upgrade Guide* – This guide details the system requirements for installing and operating vRanger, and describes how to install and upgrade the vRanger application and components.

- *Dell™ vRanger™ Integration Guide for EMC® Data Domain® Boost™* – This guide describes how to configure vRanger to work with a Data Domain Boost repository to achieve source-side deduplication of backup data.

- *Dell™ vRanger™ Integration Guide for Dell™ DR Series Disk Backup Appliance*– This guide describes how to configure vRanger to work with a Dell DR series appliance to achieve deduplication of backup data.

- *Dell™ vRanger™ Integration Guide for Dell™ NetVault™ SmartDisk - vRanger™ Edition* – This guide describes how to configure vRanger to work with a NetVault SmartDisk -vRanger Edition repository to achieve deduplication of backup data.

- *Total Virtual Data Protection - Protecting virtualized applications with Dell™ vRanger™ and Dell Recovery Manager™* – This guide describes how vRanger can work with Dell Recovery Manager products to protect key services such as Active Directory, Exchange, and SharePoint.

vRanger overview

- Virtualization overview
- vRanger overview
- Installation overview
- Configuration overview

Virtualization overview

As companies grow, their technology needs change. These changes are often implemented within complex systems running business-critical applications. Usually there is an increased demand for hardware and software resources. To manage this demand, many companies establish virtual environments. Doing so can increase an organization's agility and efficiency while lowering its costs.

Most companies have a number of specialized physical servers and workstations that are under used. Virtualizing such an environment increases and balances utilization by consolidating the physical machines into a single physical host that runs multiple virtual machines (VMs).

vRanger overview

vRanger is a backup, recovery, and backup management solution for virtual and physical environments. Featuring a new platform technology, vRanger reduces the backup window, provides smarter backup options, and offers more scalability features while using fewer resources.

vRanger capabilities:

- Performs incremental, differential or full image backups of virtual and physical machines.
- Quickly restores the entire machine or just specific files.
- Manages disaster recovery strategies and protects critical data in virtual and physical environments.
- Replicates VMware virtual machines to more than one destination, letting users accomplish multiple HA and DR objectives based on their specific needs.

What is vRanger?

vRanger is the next step in the evolution of the leading backup solution designed for the virtual infrastructure as well as physical machines. When protecting a VMware VM, vRanger captures the complete VM image (including the OS, patches, and applications) and transfers the archive (full, incremental, or differential) to a configured data repository. These image-level backups can be restored in a matter of minutes rather than the hours required by legacy file-level solutions.

vRanger was originally designed to integrate with VMware Virtual Infrastructure and vSphere at the API level, rather than through complicated scripting, and now also protects physical and Hyper-V virtual machines.

For VMware servers, vRanger uses simplified vCenter and ESX Server integration to enumerate and display your virtual environment the way you are accustomed to seeing it. vRanger also tracks and adapts to vMotion activity to avoid disruption to scheduled jobs.

To protect Hyper-V VMs, vRanger deploys an agent to the cluster or host, from which all backup functions operate. Using the VSS writers found on every Hyper-V VM, the vRanger agent captures the complete VM image (including the OS, patches, and applications) and transfers the archive (full, incremental, or differential) to a configured data repository. The vRanger Migration function is comparable to VMware vMotion in that it lets you continue protecting a server even as it moves from host to host within a cluster.

Protecting a physical machine is possible with the vRanger physical backup client. The client is installed on the physical machine you want to protect and communicates with the vRanger server, gathering information on the source server configuration, executing backup functions, and transmitting those backups to the repository through "direct-to-target" architecture.

> ⓘ **IMPORTANT:** Physical backups are only supported on Windows machines matching the list of supported operating systems. Refer to "*System requirements and compatibility*" in the *Dell vRanger Pro Installation/Upgrade Guide*.

How does it help?

Simply put, vRanger provides better data protection using fewer resources. vRanger's image-level backups for VMware VMs allow for simple, portable archives that can be restored with a few clicks of the mouse. The improved File Level Recovery (FLR) engine lets you browse archives and recover files from archives instantly. The vRanger user interface lets you easily manage the protection of all VMware, Hyper-V, and physical servers from a central location.

The intelligent Resource Manager maximizes throughput and protects infrastructure performance, while incremental backups conserve storage space. With the vRanger direct-to-target architecture, expensive proxy servers are no longer required.

Installation overview

The installation of vRanger is a simple process requiring no additional components or configuration. vRanger can be installed on either a physical or virtual machine meeting the requirements specified in the *Dell vRanger Installation and Upgrade Guide*.

Before installing

vRanger can be installed in a variety of configurations depending on your environment. Some common configuration options include:

- Database Installation
- Backup Architecture:
    - VA-based HotAdd
    - Machine-based HotAdd
    - VA-based LAN
    - Machine-based LAN
    - Machine-based SAN
- Virtual Appliance deployment

Before installing vRanger, review each of your architectural options, and decide on the best fit for your environment.

For more information on installing vRanger, refer to the *Dell vRanger Pro Installation/Upgrade Guide*.

Downloading an evaluation copy of vRanger

Evaluation versions of Dell products, including vRanger, are freely available from the Dell website. Evaluation versions are fully applications limited only by licensing.

To download an evaluation copy of vRanger

1   From an internet connected PC, access http://software.dell.com/products/vranger/.

2   Click **Download Free Trial**. Register or log in as appropriate.

3   Accept the License Agreement. Click **Submit**.

4   Select **vRanger x.x** and save the file to the preferred location.

Installing vRanger

The installation of vRanger has several options. Unless you have a valid reason not to, Dell recommends to accept the default configuration wherever possible. Complete the steps in the following procedure to install vRanger.

To install vRanger

1   Launch the installer and accept the vRanger license. The install process does not continue until the license is accepted.

2   Choose an installation directory.

3   Enter the credentials under which the vRanger services should run. The credentials used need to have local administrator privileges on the vRanger machine.

4   Select the vRanger database. You may choose to install vRanger with a new instance of SQL Server Express or on an existing SQL Server. You also need to configure DB credentials at this time.

5   The installation completes.

For more information on installing vRanger, see the *Dell vRanger Pro Installation/Upgrade Guide*.

Configuration overview

vRanger requires some basic configuration before data protection can begin. The bulk of this configuration is driven by the Startup Wizard which launches the first time the application is opened. For more information about configuring vRanger, see Configuring vRanger.

The primary configurations you need to make are described below.

Add inventory for protection

Before you can begin backups, you must add at least one of the following to the vRanger inventory:

*   one or more vCenters
*   one or more Hyper-V clusters
*   one or more vCloud Director servers
*   one or more physical servers
*   individual Hyper-V or ESXi hosts not associated with a cluster or vCenter.

To add a vCenter, you need to have credentials with administrator access to the vCenter, along with root-level credentials for each host managed by the vCenter. To add a Hyper-V cluster, you must have domain

administrator privileges. You have the option to exclude hosts from the vRanger inventory, which also exclude them from licensing.

Add repositories

Repositories are where vRanger stores the savepoints created by each backup job. You can create a repository from a standard FTP, NFS, CIFS, or SFTP share. You may also deduplicate vRanger backups by sending them to a Dell DR RDA, Data Domain DD Boost, or NetVault SmartDisk repository. To add a repository, you need the name of the server and share, as well as an account with access to that share.

Deploy virtual appliances

vRanger uses a virtual appliance for replication to/from ESXi servers, for Linux FLR, and optionally for backups and restores. Before fully utilizing vRanger, you need to deploy and configure virtual appliances to your ESXi servers. For more information, see the *Dell vRanger Pro Installation/Upgrade Guide*.

Add physical servers to inventory

vRanger includes the ability to back up and restore physical servers. Before you can backup a physical server, you must add it to the vRanger inventory and supply credentials with Administrator access to the server

Schedule a backup job

With a vCenter and a Repository configured, the next step is to schedule a backup job. To schedule a backup job, follow the sequence below.

To schedule a backup job

1   Select an object to back up. You may select any vCenter object (VM, Host, Datacenter, and so on) or a physical server.

2   To start the Backup Wizard, do one of the following:

  - Click the **Add** drop-down menu, and then click **Backup Job**.

  - Right-click the object you want to back up, and then click **Backup [selected object]**.

3   Select the disks to include or skip.

4   Select the Repository to use.

5   Select your transport method (automatic or custom).

6   Select your backup options, such as compression or enabling guest quiescing.

7   Select your backup type (Incremental, Differential, or Full) and configure Retention Policies.

8   Schedule the backup job or set to run On Demand.

For more information about the Backup Wizard or available options, see Backup.

Schedule a replication job

With a vCenter and a virtual appliance or appliances (if necessary) configured, the next step is to configure a replication job. To schedule a replication job, follow the sequence below:

To schedule a replication job

1   Select a VM to replicate.

2   To start the Replication Job Wizard, do one of the following:

   - Click the **Add** drop-down menu, and then click **Replication Job.**
   - Right-click the VM you want to replicate, and then click **Replicate [VM Name].**

3   Enter a name and description for the job.

4   Select a target host.

5   Select the type of replication: Differential or Hybrid

   Differential is the most common option.

6   Map networks for the target VM.

7   Select the disk(s) to replicate.

8   Configure job options.

9   Configure the schedule for the replication job.

10 Configure email notifications.

For more information about the Replication Wizard or available options, see Replicate.

Configuring vRanger

- Configuration overview
- Configuring vRanger through the Startup Wizard
- Manually configuring vRanger
- Supplemental instructions – additional repository types
- Understanding the vRanger virtual appliance

Configuration overview

Before you begin configuring vRanger, confirm that you have access to all user names, passwords, and root passwords for the ESX(i) hosts that you intend to use.

vRanger requires some basic configurations to ensure that the software functions properly. You complete the initial configuration through the Startup Wizard, which launches when you first install vRanger. If you would rather configure settings through the menu commands and icons on the main vRanger screen, you can do so. In either case, you can revise the settings at any time.

After completing the basic configuration process, you can use either the Startup Wizard or the Configuration Options dialog to manage your settings. Although the Configuration Options dialog offers some functionality that the Startup Wizard does not, you can use either method to manage host licensing or email server data. Only the Configuration Options dialog lets you manage resources on the task level. Given the role that tasks play in making backup and restore jobs, this resource is vital.

Configuring vRanger through the Startup Wizard

You must have an email SMTP server set up to enable notification during job creation. Later in this process, on the Email Notification Selection page, you have the option of choosing recipients. A plain text email is sent after the job completes. The notification includes the job results, if you select that option.

To configure vRanger though the Startup Wizard, complete the following tasks and processes:

1 Launching the Startup Wizard

2 Adding and selecting a license

3 Process of adding vCenter and host information

4 Adding Hyper-V System Center VMM, Cluster, and host information

5 Selecting hosts for licenses

6 Adding a physical machine

7 Selecting physical machines for licenses

8 Deploying and configuring the virtual appliance from the Startup Wizard

9 Adding a repository

Launching the Startup Wizard

Complete the following steps to launch open the Startup Wizard and begin configuring vRanger.

To launch the Startup Wizard

1   Do one of the following:

- If you have just installed the software, the Startup Wizard should automatically launch.

- If you have not, click the **Tools** drop-down menu, and then click **Startup Wizard**.

    The Startup Wizard launches.

2   Click **Next**.

Adding and selecting a license

On the Licenses page of the wizard, complete the following steps to select add and select a vRanger license. This procedure applies to adding a permanent license as well as a trial license.

To add and select a license

1   Click **Add New License**.

2   Select the license.

⚠ | **WARNING: Adding a license with features or CPUs less than the current license could result in the loss of functionality, or in some hosts or machines becoming unlicensed.**

ⓘ | **NOTE:** When applying an extended license over an existing or trial license, the highest license count is used.

3   Click **Open**.

4   Click **Save**.

5   Click **Next**.

Under Currently Installed Licenses, you see the type of license installed. If it is a trial license, it reads, License Type: Extended Trial and lists the number of days remaining in your extended trial period. Trial licenses applied display the word Trial next to the License Allowed and License Used count.

Removing an extended trial license

The extended trial license is automatically removed when the trial period has elapsed, but you may also remove it prior to the expiration of the trial period.

To remove an extended trial license

1   In the main interface, click the **Tools** menu, and then click **Options**.

2   On the Configuration Options dialog box, under the Licensing node, click **Product**.

3   On the Product License page, click **Remove Trial Extension**.

ⓘ | **NOTE:** The Remove Trial Extension button is available only when trial extension is present. When the trial extension is removed, the button is also removed.

4   In the Licensing window, click **Yes** to confirm.

5   Under **Currently Installed Licenses**, confirm that the primary license is being applied.

6   In the Virtual Machine Hosts and Physical Machines nodes, confirm your license allocation.

Process of adding vCenter and host information

The VirtualCenter and Host Information page displays. Before vRanger can add source objects to the My Inventory pane, you must configure a connection to a vCenter and/or an ESX(i) Server.

To add vCenter and host information, complete the following tasks as necessary:

1   Adding a vCenter. Complete this task to add a vCenter and all associated hosts.

2   Entering host credentials. Complete this task to add credentials to hosts that appeared after adding a vCenter.

3   Adding hosts. Complete this task to add individual hosts not associated with a vCenter.

Adding a vCenter

Complete the following steps to add a vCenter to your vRanger configuration.

To add a vCenter

1   In the VirtualCenters section, click **Add**.

    The Add VirtualCenter Credentials dialog box appears.

2   In the **DNS Name or IP** text box, enter the FQDN or IP address of the vCenter server.

3   In the **User Name** text box, enter the user name of an account with privileges on the vCenter server. For the required permissions for a vRanger vCenter account, see "Configuring vCenter permissions" in the "Before You Install" chapter of the *vRanger Installation and Upgrade Guide*.

> (i) | **NOTE:** The user name for the vCenter credential should be entered in the "username@domain" format, rather than "domain\username". In some cases, the domain may not be required. Avoid special characters in the user name.
>
> If these credentials are changed in the future, you will need to restart the vRanger Service to recognize the changes.

4   In the **User Password** text box, enter the password for the account used above.

> (i) | **NOTE:** Avoid special characters in the password.

5   In the **Port Number** text box, enter the port to be used for communication. The default port is 443.

6   Click **Connect**.

    The dialog box closes and the vCenter displays in the VirtualCenters section and on the **VirtualCenter and Host Information** page. The hosts managed by that vCenter display in the Hosts section.

    Note that the hosts displayed show the icon . There are four key indicators shown in the icon:

    • The large gold key indicates that the host has been issued a vRanger license.

    • The green dot indicates that the host has been assigned a backup license.

    • The blue dot indicates that the host has been assigned a replication license.

    • The authentication method for the host is indicated by the icon superimposed on the host icon:

       • If the host is authenticated with vCenter credentials only, the vCenter icon appears superimposed over the host icon: .

- If the host is authenticated with host credentials, a gold key is superimposed over the host icon: .

⚠ **WARNING: vCenter credentials are sufficient for operations that use only the vStorage API. You need to apply credentials to each host for backup and replication operations to use the Service Console.**

Entering host credentials

Complete the following steps to enter the credentials required to connect to the host or hosts.

To enter host credentials

1 In the **Hosts** section, select one or more hosts. To select all hosts, select the upper-most check box.

2 Click **Edit**.

   The **Host Connections** dialog appears. If you have selected multiple hosts, the DNS Name or IP section displays "[multiple connections]"

3 In the **User Name** text box, enter an account for the host. You may use the root account if connections through the root are enabled, or a non-root account. If the account does not yet exist on the host, select **Add user if user does not exist**.

4 In the **User Password** text box, enter the password for the account used above.

5 In the **Root Password** text box, enter the root password for the selected hosts. If you hosts use different root passwords, you need to configure them in like groups based on the root password.

6 In the **Port Number** text box, enter the port to be used for communication. The default port is 22.

7 Click **Connect**. The icon next to the host changes to , showing that the vCenter icon has been replaced with a key.

   If there is an authentication problem, such as connections through the root account being disabled, the status icon displays an alarm:

   If this happens, you need to edit the host connection to use a non-root user account. Make sure that the Add user if user does not exist option is selected.

   ⓘ **NOTE: ESXi requires the use of the root account. You cannot use a non-root account with ESXi.**

8 When all connections have been made and verified, click **Next**.

Adding hosts

If you have hosts that are not managed by a vCenter, you can add them individually.

To add hosts

1 In the Host section, click **Add**.

2 In the **DNS Name or IP** text box, enter the FQDN or IP address of the Host.

3 In the **User Name** text box, enter an account for the host. You may use the root account if connections through the root are enabled, or a non-root account. If the account does not yet exist on the host, select **Add user if user does not exist**.

4 In the **User Password** text box, enter the password for the account used above.

5 In the **Root Password** text box, enter the password for the root account.

6 In the **Port Number** text box, enter the port to be used for communication. The default port is 22.

7 Click **Connect**. The Host displays in the Hosts section.

8 To continue in the wizard, click **Next**.

Adding Hyper-V System Center VMM, Cluster, and host information

The Hyper-V Cluster and Host Information page of the Startup Wizard displays. Before vRanger can add source objects to the My Inventory pane, you must configure a connection to a Hyper-V System Center VMM, cluster, or host.

To add a Hyper-V System Center VMM, cluster or host, complete the following tasks as necessary:

1   Adding a Hyper-V System Center Virtual Machine Manager

2   Adding a Hyper-V Cluster

3   Adding hosts

Adding a Hyper-V System Center Virtual Machine Manager

Complete the steps in the following procedure to add a Hyper-V System Center VMM.

> ⓘ   **IMPORTANT:** Before adding the System Center VMM connection, you must install the System Center VMM Administrative Console on the vRanger machine.

To add a Hyper-V System Center VMM

1   In the Hyper-V System Center Virtual Machine Managers section, click **Add**.

2   In the **DNS Name or IP** text box, enter the FQDN or IP address of the Hyper-V System Center VMM.

3   In the **User Name** text box, enter the user name of an account with domain administrator privileges on the System Center VMM.

4   In the **User Password** text box, enter the password for the account used above.

5   If this is a new System Center VMM, or if you have removed the vRanger Hyper-V agent, select **Install agent on host.**

6   Configure the ports as follows:

   • In the **Agent Port Number** text box, enter the port you want vRanger to use to communicate with the vRanger agent installed on each Hyper-V host. This port must be open between vRanger and each Hyper-V server.The default port is 8081.

   • In the **SCVMM Port Number** text box, enter the port you want vRanger to use to communicate with the System Center VMM server. The default port number is 8100.Click **Connect**.

   > ⓘ   **IMPORTANT:** This port is configured during System Center VMM installation. If you chose a port number different from the default value, enter that value here.

7   Click **Next**.

   The Hyper-V System Center VMM displays in the Hyper-V System Center Virtual Machine Managers section. The Hosts managed by that System Center Virtual Machine Manager displays in the Hosts section.

Adding a Hyper-V Cluster

Complete the steps in the following procedure to add a Hyper-V cluster.

To add a Hyper-V cluster

1   In the Hyper-V Cluster section, click **Add**.

2   In the **DNS Name or IP** text box, enter the FQDN or IP address of the Hyper-V cluster.

3   In the **User Name** text box, enter the user name of an account with domain administrator privileges on the cluster.

4   In the **User Password** text box, enter the password for the account used above.

5   If this is a new cluster, or if you have removed the vRanger Hyper-V agent, select **Install agent on host.**

6   In the Port Number text box, enter the preferred port you want vRanger to use to communicate with the Hyper-V cluster on the source server. This port must be open between vRanger and each Hyper-V server. The default port number is 8081.

7   Click **Connect**.

8   Click **Next**.

The Hyper-V cluster displays in the Hyper-V Clusters section. The Hosts managed by that cluster displays in the Hosts section.

Adding hosts

If you have hosts that are not part of a cluster, you can add them individually.

To add hosts

1   In the Host section, click **Add**.

2   In the **DNS Name or IP** text box, enter the FQDN or IP address of the Host.

3   In the **User Name** text box, enter an account for the host.

4   In the **User Password** text, enter the password for the account used above.

5   If this is a new host, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.

6   In the Port Number text box, enter the preferred port you want vRanger to use to communicate with the Hyper-V host on the source server. This port must be open between vRanger and each Hyper-V server. The default port number is 8081.

7   Click **Connect**.

The Host displays in the Hosts section.

Selecting hosts for licenses

The Host Licensing Information page displays. It displays total licensed sockets, used sockets, and the remaining days left in a trial period.

ⓘ | **NOTE:** VM licenses include all virtual machines. Licenses for both VMware and Hyper-V VMs are drawn from the same socket count.

To select hosts for licenses

1   View the **CPUs Allowed** value to confirm the number of sockets that are licensed for protection. The number of licenses currently allocated is shown in the **CPUs Used** column.

The Licensed Hosts section lists the managed hosts, the number of sockets for each host, and the licensed features for that host.

The CPUs column indicates the license "cost" of each host.

2   In the Pro Backup or Replication column, select or deselect hosts until you have used all of the available licenses, or until all of your hosts are protected.

3   Click **Save**.

4   Click **Next**.

Adding a physical machine

The Physical Machine Information page displays. Before vRanger can add physical source objects to the My Inventory pane, you must configure a connection to a physical server.

To add a physical machine

1   In the Physical Machines section, click **Add**.

2   n the **DNS Name or IP** text box, enter the FQDN or IP address of the server.

3   In the **User Name** text box, enter an account for the server.

4   In the **User Password** text box, enter the password for the account used above.

5   If this is a new server, or if you have removed the vRanger agent, select **Install agent on machine**.

In the Agent Location text box, enter the preferred directory (on the physical machine) to which the physical client should be installed. The default installation location is:

> *C:\Program Files\Dell\vRangerPhysicalClient*

6   In the Port Number text box, enter the preferred port for vRanger to use to communicate with the physical client on the source server. This port must be open between vRanger and each physical server. The default port number is 51000

7   Click **Connect**.

The server displays in the My Inventory pane. You may also create a Backup Group to combine multiple physical servers into one backup job. See the section Adding a custom backup group for more information.

Manually installing the physical client

vRanger requires a physical backup client to be installed on the source physical server. This is normally done when the source server is added to the vRanger inventory, but in some circumstances it may be necessary to install the client manually.

To manually install the physical client

1   On the vRanger server, find the physical client installation file - *vRangerPhysicalClient.exe*. The file is located in the "Service" sub-directory of the vRanger installation folder. By default this is:

> *C:\Program Files\Dell\vRanger\Service*

2   Copy the client to the preferred location on the source server

3   From the source server's command line, run the vRangerPhysicalClient.exe command modified by the parameters below:

- **USERNAME** – This is required information. This is the username for an account with administrative privileges Log on as a service" rights on the source server.

- **PASSWORD** – This is required information. This is the password for the account used above.

- **DOMAIN** – This is optional. If no value is entered, the account used above is considered a local account. If you use a network account, enter the domain for the account here.

- **INSTALL_DIR** – This is optional. Default install directories are:

> *C:\Program Files\Dell\vRangerPhysicalClient*

- **PORT** – This is optional. The default port is 51000

Command examples

The examples below show the proper syntax for using the installer.

- To accept the defaults, you must only provide credentials for the source server:

```
vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
```

- You may also use a network account:

```
vRangerPhysicalClient.exe /DOMAIN="development "/USERNAME="user1"
    /PASSWORD="Password"
```

- To change the install location or port number, add that parameter with the preferred value:

```
vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
    /INSTALL_DIR="c:\vRangerPhysicalClient"
```

```
vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
    /PORT=52000
```

Selecting physical machines for licenses

Before you start this process, make sure that you have imported the license file and confirmed that it is adequate for managing the number of physical machines you want to protect. The Currently Installed Licenses section includes the total number of physical machine licenses available, and the number being used currently. These numbers increment and decrement based on the selections you make on this page.

To select physical machines for licenses

1   In the **Currently Installed Licenses** section, view the **Machines Allowed** value to confirm the total number of physical machines licenses available The number of licenses currently allocated is shown in the **Machines Used** column.

2   The **Licensed Physical Machine** section lists the physical machines in vRanger's inventory, and indicates whether each machine is licensed for backup

    a   In the **Physical Backup** column, select or deselect machines until you have used all of the available licenses, or until all of your machines are protected.

3   Click **OK**.

Deploying and configuring the virtual appliance from the Startup Wizard

vRanger uses a virtual appliance (VA) for both Linux file-level recovery and for replication to and from VMware ESXi servers.

There are two ways to deploy and configure a VA: the Startup Wizard and the Tools menu.

If you do not want to complete the Virtual Appliance Information page of the Startup Wizard now, you may skip this step and continue with the Startup Wizard. You can access the Virtual Appliance Configuration dialog at any time by way of the Options available from the Tools drop-down menu.

To complete the Virtual Appliance Information page of the Startup Wizard and deploy and configure a VA now, complete the following procedure.

To deploy and configure the virtual appliance from the Startup Wizard

1   On the Virtual Appliance Information page of the Startup Wizard, launch the Virtual Appliance Deployment Wizard from the Startup Wizard by clicking **Deploy Virtual Appliance**.

2   Complete the deployment wizard by following the steps in Deploying the virtual appliance by using the Virtual Appliance Deployment Wizard.

3   To add a new virtual appliance configuration, click **Add**, and then complete the following steps:

    a   In the Add Virtual Appliance Configuration dialog, select a virtual appliance from the inventory tree.

b Under Virtual Appliance Properties, you do any of the following:

- Select **Override IP Address**, and then enter a new IP address in the **IP Address** text box.

- Enter a root password for the VA in the **Root Password** text box.

- Select **Use as default virtual appliance for cluster**, to use this VA for all machines that are a part of the associated cluster.

c Click **OK**.

4 To configure an existing virtual appliance, select a VA from the list, and then click **Edit**.

In the Modify Virtual Appliance Configuration dialog, you can edit any of the following settings:

- Virtual Appliance Properties
- Virtual Appliance Options
- Replication
- Scratch Disk
- Password
- Linux File Level Restore

For more information, see Changing the virtual appliance configuration.

5 To delete a virtual appliance, do the following:

a Select a VA from the list.

b Click **Remove**.

The Removing VA dialog appears.

c Select the job you want to remove.

d Click **OK**.

e In the Confirm Delete dialog, click **OK**.

If you want remove the entire VA rather than a single job, first select **Delete the virtual appliance from the host**, and then click **OK**.

6 Under Linux FLR Virtual Appliance, if you want to plan for Linux File Level Recovery, select the virtual appliance you want to use from the drop-down list.

7 Click **Next**.

Adding a repository

vRanger uses repositories to store backup archives. Repositories can be one of the following types:

- CIFS
- NFS (version 3)
- FTP
- SFTP
- **NetVault SmartDisk** - Dell Software's disk-based data-deduplication option which reduces storage costs with byte-level, variable-block-based software deduplication. For more information on NetVault SmartDisk, see http://software.dell.com/products/netvault-smartdisk/ or the *Dell vRanger Integration Guide for NetVault SmartDisk*.
- **Dell Rapid Data Access (RDA)** - Provided by the Dell DR Series appliances - purpose-built, disk backup appliances that use Dell deduplication technology to significantly improve backup and recovery processes. For more information on Dell DR Series appliances, see http://software.dell.com/products/dr-series-disk-backup-appliances/ or the *Dell vRanger Integration Guide for Dell DR Series Disk Backup Appliance*.

- **EMC Data Domain Boost (DDB)** - Integrating EMC® Data Domain® Boost™ with vRanger is achieved by adding a Data Domain appliance running DD Boost to vRanger as a repository. Backups written to that repository are deduplicated according to your configuration. For more information on EMC Data Domain Boost, see http://www.emc.com/data-protection/data-domain/data-domain-boost.htm or the *Dell vRanger Integration Guide for EMC® Data Domain® Boost™*.

ⓘ **NOTE:** Dell RDA and EMC Data Domain Boost repositories support repository replication. For more information, refer to Managing Repository Replication.

The procedure below shows mounting a CIFS share to the My Repositories pane. The procedures for the other repository types are described in Supplemental instructions – additional repository types

To add a repository

1   Under Repositories, select **Windows Share (CIFS)**.

2   Populate the **Repository Name** text box.

   This value displays in the **My Repositories** pane.

3   Populate the **Description** text box.

4   Enter a **username** and **password** in the appropriate text boxes.

5   Select a **Security Protocol** from the drop-down: NTLM (Default), or NTLM v2

6   In the **Server** text box, type the UNC path to the preferred repository directory. Alternatively, you may enter a partial path and click **Browse** to find the target directory.

   ⓘ **NOTE:** You must enter a valid username and password before using the browse functionality.

⚠ **WARNING: If you want to use the Encrypt all backups feature, make certain to retain the password you enter in the following steps. There is no back-door or admin-level password. If the password is unknown, the backups are not usable.**

7   Select **Encrypt all backups to this repository** if you want these backups to be password-protected.

   ⓘ **NOTE:** Encryption is not supported for NetVault SmartDisk and Data Domain Boost repositories.

8   Enter a Password for the encrypted repository - confirm the password by re-entering it.

9   Click **Save** — the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

10 vRanger checks the configured repository location for existing manifest data to identify existing savepoints. If any are found, you are prompted to take one of three actions:

- **Import as Read-Only** – With this option, all savepoint data is imported into the vRanger database, but only for restores. You cannot back up to this repository.

- **Import** – All savepoint data is imported into the vRanger database. vRanger can use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – The savepoint data is retained on the disk, but cannot be imported into vRanger. vRanger ignores the existing savepoint data and treats the repository as new.

11 Click **Next**.

Setting up the SMTP server

On the Mail Server Information page, you can configure an SMTP server, Select one that does not require authentication so that there are no obstacle to email notification transmission.

To set up the SMTP server

1   In the **SMTP Server** text box, enter the DNS name or IP of an SMTP server.

2   In the **Port** text box, enter the port to be used for communication. The default port is 25.

3   In the **From Address** text box, enter the email address that vRanger should use.

> ⓘ | **NOTE:** Make sure that you use a valid email account registered to the SMTP server that you added.

4   Click **Send Test Email**.

5   If preferred, select **Do not send email notifications for successful jobs**.

Adding an email address

Complete the following steps on the Mail Server Information page of the wizard to add an email address to the vRanger address book.

To add an email address

1   To add recipients, click **Add**.

2   Populate the **Name** and **Email Address** text boxes.

3   Click **OK** to add the information to the Email Addresses list.

4   Repeat steps 1-3 to include additional contacts.

5   Click **Finish**.

Manually configuring vRanger

The following instructions describe how to configure vRanger without using the Startup Wizard by manually adding the servers you want to protect. You can manually add VMware and Hyper-V virtual machines, physical machines, and repositories, as well as the configurations necessary for email notifications.

The process of manually configuring vRanger includes the following tasks:

- Adding a VirtualCenter
- Adding a vCloud Director
- Adding an ESX server
- Adding a Hyper-V System Center VMM, Cluster, or host
- Adding a physical server
- Adding a repository
- Configuring an SMTP server
- Configuring display language
- Configuring a virtual appliance
- Importing a license file
- Managing virtual machine host licensing
- Managing physical machine licensing
- Managing job configuration settings
- Enabling cataloging options
- Managing Transport Failover Options

Adding a VirtualCenter

Complete the steps in the following procedure to manually add a VMware VirtualCenter (vCenter) for vRanger backup.

To add a VirtualCenter

1 On the **My Inventory** view, click the **VMware** tab.

2 Click the **Add** icon, and then click **Virtual Center**.

The Add VirtualCenter Credentials dialog displays.

3 In the **DNS Name or IP** text box, enter the FQDN or IP address of the vCenter server.

4 In the **User Name** text box, enter the user name of an account with privileges on the vCenter server. See Configuring vCenter permissions for the required permissions for a vRanger vCenter account.

> ⓘ **NOTE:** The user name for the vCenter credential should be entered in the "username@domain" format, rather than "domain\username". In some cases, the domain name may not be required. Avoid special characters in the user name.
>
> If these credentials are changed in the future, you will need to restart the vRanger Service to recognize the changes.

5 In the **User Password** text box, enter the password for the account used above.

> ⓘ **NOTE:** Avoid special characters in the password.

6 Click **Connect**.

vRanger attempts to connect to the vCenter using the current configuration. If it is successful, the vCenter displays in the **My Inventory** pane.

Adding a vCloud Director

Complete the steps in the following procedure to manually add a VMware vCloud Director server for vRanger backup.

To add a vCloud Director

1 On the **My Inventory** view, click the **vCloud** tab.

2 Click the **Add** icon, and then click **vCloud Director**.

The **Add vCloud Director Credentials** dialog displays.

3 In the **DNS Name or IP** text box, enter the FQDN or IP address of the vCloud Director server.

4 In the **User Name** text box, enter the user name of an account with administrator privileges on the vCloud Director server.

5 In the **User Password** text box, enter the password for the account used above.

> ⓘ **NOTE:** Avoid special characters in the password.

6 Click **Connect**.

vRanger attempts to connect to the vCloud Director server using the current configuration. If it is successful, the vCloud Director displays in the **My Inventory** pane.

You will be prompted to add credentials for each vCenter server attached to the vCloud Director.

7 In the Add vCloud Director Credentials dialog, identify any vCenter servers for which credentials are not configured. The authentication status is indicated by a connection icon in the left most column. A red disconnected icon indicates that a vCenter needs authentication.

8 Add the user name and password for the vCenter servers as required. Click **Connect** to authenticate.

9 Click **Done**.

Adding an ESX server

Complete the steps in the following procedure to manually add a VMware ESX server to vRanger protection.

To add an ESX server

1 On the **My Inventory** view, click the **VMware** tab.

2 Click the **Add** icon, and then click **ESX Host**.

> ⓘ | **NOTE:** If the connection is unsuccessful, confirm that the configuration settings are correct and that the vCenter is visible to the computer where vRanger is installed.

3 In the **DNS Name or IP** text box, enter the FQDN or IP address of the Host.

4 In the **User Name** text box, enter an account for the host. You may use the root account if connections through the root are enabled, or a non-root account. If the account does not yet exist on the host, select **Add user if user does not exist**.

5 In the **User Password** text box, enter the password for the account used above.

6 In the **Root Password** text box, enter the password for the root account.

7 Click **Connect**.

The host is added to the tree structure under the vCenter. Confirm this noting that a gold key icon displays adjacent to it in the My Inventory pane.

Adding a Hyper-V System Center VMM, Cluster, or host

Complete the steps in the following procedure to manually add a Hyper-V System Center VMM, cluster, or host to vRanger protection.

To add a Hyper-V Host

1 On the **My Inventory** view, click the **Hyper-V** tab.

2 Click the **Add** icon, and then click **Hyper-V Host**

The Add Credentials dialog displays.

3 In the **DNS Name or IP** text box, enter the FQDN or IP address of the Host.

4 In the **User Name** text box, enter an account for the host.

5 In the **User Password** text, enter the password for the account used above.

6 If this is a new host, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.

In the Port Number text box, enter the preferred port you want vRanger to use to communicate with the Hyper-V host on the source server. This port must be open between vRanger and each Hyper-V server. The default port number is 8081.

To add a Hyper-V System Center VMM

ⓘ | **IMPORTANT:** Before adding the System Center VMM connection, you must install the System Center VMM Administrative Console on the vRanger machine.

1   On the **My Inventory** view, click the **Hyper-V** tab.

1   Click the **Add** icon, and then click **Hyper-V System Center VMM.**

2   In the **DNS Name or IP** text box, enter the FQDN or IP address of the Hyper-V System Center VMM.

3   In the **User Name** text box, enter the user name of an account with domain administrator privileges on the System Center VMM.

4   In the **User Password** text box, enter the password for the account used above.

5   If this is a new System Center VMM, or if you have removed the vRanger Hyper-V agent, select **Install agent on host.**

6   Configure the ports as follows:

   - In the **Agent Port Number** text box, enter the port you want vRanger to use to communicate with the vRanger agent installed on each Hyper-V host. This port must be open between vRanger and each Hyper-V server.The default port is 8081.

   - In the **SCVMM Port Number** text box, enter the port you want vRanger to use to communicate with the System Center VMM server. The default port number is 8100.Click **Connect.**

     ⓘ | **IMPORTANT:** This port is configured during System Center VMM installation. If you chose a port number different from the default value, enter that value here.

7   Click **Next.**

The Hyper-V System Center VMM displays in the Hyper-V System Center Virtual Machine Managers section. The Hosts managed by that System Center Virtual Machine Manager displays in the Hosts section.

Adding a Hyper-V Cluster

Complete the steps in the following procedure to add a Hyper-V cluster.

To add a Hyper-V cluster

1   On the **My Inventory** view, click the **Hyper-V** tab.

2   Click the **Add** icon, and then click **Hyper-V Failover Cluster.**

3   In the **DNS Name or IP** text box, enter the FQDN or IP address of the Hyper-V cluster.

4   In the **User Name** text box, enter the user name of an account with domain administrator privileges on the cluster.

5   In the **User Password** text box, enter the password for the account used above.

6   If this is a new cluster, or if you have removed the vRanger Hyper-V agent, select **Install agent on host.**

7   In the Port Number text box, enter the preferred port you want vRanger to use to communicate with the Hyper-V cluster on the source server. This port must be open between vRanger and each Hyper-V server. The default port number is 8081.

8   Click **Connect.**

9   Click **Next.**

The Hyper-V cluster displays in the Hyper-V Clusters section. The Hosts managed by that cluster displays in the Hosts section.

Adding hosts

If you have hosts that are not part of a cluster, you can add them individually.

To add hosts

1. In the Host section, click **Add**.

2. In the **DNS Name or IP** text box, enter the FQDN or IP address of the Host.

3. In the **User Name** text box, enter an account for the host.

4. In the **User Password** text, enter the password for the account used above.

5. If this is a new host, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.

6. In the Port Number text box, enter the preferred port you want vRanger to use to communicate with the Hyper-V host on the source server. This port must be open between vRanger and each Hyper-V server. The default port number is 8081.

7. Click **Connect**.

    The Host displays in the Hosts section.

Adding a physical server

Complete the steps in the following procedure to manually add a physical server to vRanger protection.

To add a physical server

1. From the **My Inventory** view, click the **Physical** tab.

2. Click the **Add** icon, and then click **Physical Machine**.

    The Add Physical Machine Credentials displays.

3. In the **DNS Name or IP** text box, enter the FQDN or IP address of the server.

4. In the **User Name** text box, enter an account for the server.

5. In the **User Password** text box, enter the password for the account used above.

6. If this is a new server, or if you have removed the vRanger agent, select **Install agent on machine.**

    In the Agent Location text box, enter the preferred directory (on the physical machine) to which the physical client should be installed. The default installation location is:

    > For 64-bit servers: *C:\Program Files\Dell\vRangerPhysicalClient*

7. In the Port Number text box, enter the preferred port for vRanger to use to communicate with the physical client on the source server. This port must be open between vRanger and each physical server. The default port number is 51000

8. Click **Connect**.

    The server displays in the My Inventory pane. You may also create a Backup Group to combine multiple physical servers into one backup job. See the section Adding a custom backup group for more information.

Manually installing the physical client

vRanger requires a physical backup client to be installed on the source physical server. This is normally done when the source server is added to the vRanger inventory, but in some circumstances it may be necessary to install the client manually.

To manually install the physical client

1   On the vRanger server, find the physical client installation file - *vRangerPhysicalClient.exe*. The file is located in the "Service" sub-directory of the vRanger installation folder. By default this is:

>   *C:\Program Files\Dell\vRanger\Service*

2   Copy the client to the preferred location on the source server

3   From the source server's command line, run the vRangerPhysicalClient.exe command modified by the parameters below:

- **USERNAME** – This is required information. This is the username for an account with administrative privileges Log on as a service" rights on the source server.

- **PASSWORD** – This is required information. This is the password for the account used above.

- **DOMAIN** – This is optional. If no value is entered, the account used above is considered a local account. If you use a network account, enter the domain for the account here.

- **INSTALL_DIR** – This is optional. Default install directories are:

>   *C:\Program Files\Dell\vRangerPhysicalClient*

- **PORT** – This is optional. The default port is 51000

Command examples

The examples below show the proper syntax for using the installer.

- To accept the defaults, you must only provide credentials for the source server:

```
vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
```

- You may also use a network account:

```
vRangerPhysicalClient.exe /DOMAIN="development "/USERNAME="user1"
    /PASSWORD="Password"
```

- To change the install location or port number, add that parameter with the preferred value:

```
vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
    /INSTALL_DIR="c:\vRangerPhysicalClient"

vRangerPhysicalClient.exe /USERNAME="administrator" /PASSWORD="Password"
    /PORT=52000
```

Adding a repository

The following instructions describe how to add a CIFS repository to vRanger. For information about adding other types of repositories, see Supplemental instructions – additional repository types.

To add a repository

1   To add repository, navigate to the My Repositories pane. Click to highlight the preferred repository type.

2   Right-click the node, select **Add**, and then click **Windows Share CIFS**.

> ⓘ | **NOTE:** These instructions describe the process for adding a CIFS repository. The procedures for the other repository types are described in Supplemental Instructions – Additional Repository Types.

3   On the **Add Repository** dialog, populate the text boxes described below:

- **Repository Name** – Enter a name for the repository.

- **Description** – Optionally, describe the repository.

- **User Name** and **User Password** text boxes – Enter the credentials for accessing the CIFS share.

- Select a **Security Protocol** from the drop-down: NTLM (Default), or NTLM v2.

- **Server** – Enter the path to the preferred directory. Alternatively, you may enter the server IP or DNS name and click **Browse**, then browse to the correct path.

⚠ WARNING: **Paths for CIFS repositories are case sensitive. When entering the path for a CIFS directory, ensure that the path is entered in the correct case.**

4   Select **Encrypt all backups to this repository** if you want these backups to be password-protected.

5   Enter a Password for the encrypted repository — confirm the password by re-entering it.

⚠ WARNING: **If you want to use the Encrypt all backups feature, make certain to retain the password you enter in the following steps. There is no back-door or admin-level password. If the password is unknown, the backups are not usable.**

6   Click **Update** — the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

7   Click **OK**.

Managing Repository Replication

vRanger supports repository replication for Dell RDA and EMC Data Domain Boost repositories. Repositories configured for replication can be synchronized in one of three ways:

- Automatically, after a successful backup task to that repository
- As a separate job on a scheduled basis
- Manually, via the Synchronize option

Configuring a repository for automatic replication

When a managed repository is configured for savepoint replication, and a backup task completes successfully, each savepoint will also be replicated to the replication repository.

To enable automatic savepoint replication:

1   From the **Tools** menu on the vRanger toolbar, select **Options**. Under the **Repositories** node, select **Replication.**

  ⓘ NOTE: You may also right-click on the target repository (in the My Repositories View) and select **Repository Replication Options**.

2   Select **Enable savepoint replication for a successful backup job task.** Click **Ok**.

Scheduling Repository Synchronization

When using scheduled repository synchronization, savepoints will be replicated to the replication repository according to a configured schedule.

To schedule a repository synchronization:

1   From the **Tools** menu on the vRanger toolbar, select **Options**. Under the **Repositories** node, select **Replication.**

2   In the **Repository Replication Configuration** pane, find the **Repository Name** column. Select the desired repository. Note that the repositories listed are the target repositories.

3   Select **Schedule repository synchronization**.

4   Configure the replication schedule as desired, using the information below as a guide.

  a   **Start**- in the drop-down, select the time for the replication task to begin.

  b   **Recurrence Pattern** - In this section, you establish how often the changes should be synchronized. There are five options within this section:

- **Daily** – The daily option can be scheduled to synchronize the repository every weekday or every *x* number of days.

- **Weekly** – Repository synchronization can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run synchronization tasks can be configured.

- **Monthly** – The monthly option offers the following configurations:

    - **Day** *[x]* **of every** *[y]* **month:**

        x can be any value from 1 to 31. This determines the day of the month on which the synchronization job occurs.

        y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.

    - **The** *[f] [d]* **of every** *[y]* **month(s)**:

        f can be either: first, second, third, fourth or last.

        d can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
        y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the synchronization task to run every other month.

- **Yearly** – The yearly option offers the following configurations:

    - **Every** *[m][x]*:

        m is any month of the year. This determines the month of the year in which the synchronization occurs.

        x can be any value from 1 to 31. This determines the day of the month on which the synchronization occurs.

    - **The** *[f][d]* **of** *[m]*:

        f can be either: first, second, third, fourth, or last.

        d can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
        m is any month of the year. This determines the month of the year in which the synchronization occurs.

- **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between backup jobs. The interval selected must be greater than or equal to five minutes.

5 Click **OK** to schedule the job.

To manually synchronize a repository:

1 From the My Repositories view, right-click on the source or replication repository and select **Synchronize**.

2 At the confirmation prompt, click **Yes**.

Disabling repository replication

To disable repository replication;

1 From the **Tools** menu on the vRanger toolbar, select **Options**. Under the **Repositories** node, select **Replication.**

OR

Right-click on the replication repository, and select Repository Replication Options.

2   Clear the check box for **Enable savepoint replication for a successful backup job task** and/or **Schedule repository synchronization.**

3   Click **OK**. Replication will be disabled as indicated by a red circle icon containing a white exclamation point.

Changing a replication repository to a primary repository

If your primary repository becomes corrupt or otherwise unavailable, you can quickly change your replication repository to a primary repository to continue backup and recovery operations.

To change a replication repository to a primary repository:

1   From the My Repositories view, right-click the replication repository and click **Remove**.

> ⓘ   IMPORTANT: Ensure that Delete all savepoints in the repository is not selected.

2   Click **OK**.

3   In the My Repository view, click Add.

4   Select the appropriate repository type (either **EMC Data Domain Boost** or **Dell Rapid Data Access (RDA)**).

5   Complete the Add Repository dialog as appropriate for your repository type. Use the information in Adding an EMC Data Domain Boost (DDB) repository or Adding a Dell Rapid Data Access (RDA) repository as required. Click **OK**.

6   vRanger will detect that the repository being added contains savepoint data, and display the **Warning: Existing Repository Found** dialog. To re-configure to a primary repository, click **Import**.

Configuring an SMTP server

Complete the steps in the following procedure to manually configure an SMTP server for email notifications from vRanger.

To configure an SMTP server

1   Click **Tools**, and then click **Options**.

2   Under the General node on the **Configuration Options** page, click **Email Configuration**.

3   Populate the **SMTP Server** text box.

4   Leave the default (25) in the **Port** text box.

5   In the **From Address** text box, enter an email address from which notifications should be sent.

6   Click **Send Test Email**.

The confirmation message confirms the server connection and successful transmission to the email address you entered.

7   Click **OK**.

8   If preferred, select **Do not send email notifications for successful jobs.**

Adding an email address

Complete the steps in the following procedure to manually add the email addresses you want to receive email notifications from vRanger.

To add an email address

1 Click **Tools**, and then click **Options**.

2 Under the General node on the **Configuration Options** page, click **Email Configuration**.

3 In the **Email Addresses** section, click **Add**.

4 Click the **Add** button.

5 Enter the contact name and email address, and click **OK**.

Editing an email address

Complete the steps in the following task to edit email addresses in vRanger.

To edit an email address

1 Click **Tools**, and then click **Options**.

2 Under the **General** node, click **Email Configuration**.

3 In the **Email Addresses** section, select the email address that you want to edit. Click **Edit**.

4 Update the information and click **Update**.

5 The information is updated and moved to the **Email Addresses** section.

6 Click **OK**.

Removing an email address

Complete the steps in the following procedure to remove an email address from vRanger.

To remove an email address

1 Click **Tools**, and then click **Options**.

2 Under the General node on the **Configuration Options** page, click **Email Configuration**.

3 In the **Email Addresses** section, select the email address that you want to delete.

4 Click **Remove**. The listing is deleted.

5 Click **OK**.

Configuring display language

vRanger lets you configure the language used when displaying the interface.

To configure display language

1 Click **Tools**, and then click **Options**.

2 Under the General node on the **Configuration Options** dialog, click **Language**.

3 From the **Available Languages** drop-down, select the language to be used. If preferred, select **based on system locale** to make an automatic setting based on the vRanger machine's locale setting.

4 Click **OK**.

Configuring a virtual appliance

vRanger uses a virtual appliance for both Linux file-level recovery and for replication to and from VMware ESXi servers. The Virtual Appliance Configuration dialog lets you add a configured VA to vRanger, and select the VA to use for Linux file-level recovery. Virtual appliances deployed using the Virtual Appliance Deployment Wizard is added to vRanger automatically.

> ⓘ | **NOTE:** You must have deployed and configured at least one virtual appliance before you can add it to vRanger. See Understanding the vRanger virtual appliance for instructions on virtual appliance deployment.

To configure a virtual appliance

1 From the **Configuration Options** dialog, select **Virtual Appliance Configuration**.

2 Click **Add**.

The **Add virtual appliance configuration** dialog appears, showing the vCenter inventory tree.

3 Browse the inventory tree, and select the virtual appliance.

4 Enter the root password for the virtual appliance, and click **OK**.

The virtual appliance appears in the Deployed Virtual Appliances section.

Importing a license file

Before you start this process, confirm that you have placed your license file in a location that you can reach easily.

To import a license file

1 Click **Tools**, and then click **Options**.

2 Under the **Licensing** node on the **Configuration Options** dialog, click **Product**.

3 Click **Add New License**.

4 Browse to the location where you stored the license file. Click to select it and click **Open**. The new license information is added to the dialog.

5 Click **Apply**.

The dialog refreshes with the new license information added.

6 Click **OK**.

Adding an extended trial license

The process of adding an extended trial license is the same as adding a permanent license. Before you start this process, confirm that you have placed your license file in a location that you can reach easily.

To add an extended trial license

1 Click **Tools**, and then click **Options**.

2 Under the **Licensing** node on the **Configuration Options** dialog, click **Product**.

3 Click **Add New License**.

4 Browse to the location where you stored the license file. Click to select it and click **Open**. The new license information is added to the dialog.

5 **Click Apply.** The dialog refreshes with the new license information added.

6 Under the **Currently Installed Licenses** heading, you see "License Type: Extended Trial", as well as the number of days remaining in your extended trial period.

Removing an extended trial license

Extended trial licenses are automatically removed when the trial period has elapsed, but you may also remove then prior to the expiration of the trial period.

To remove an extended trial license

1 In the Configuration Options > Licensing node, click **Remove Trial Extension**.

2 Under the **Currently Installed Licenses**, confirm that the primary license is being applied.

3 In the Virtual Machine Hosts and Physical Machines node, confirm your license allocation.

Managing virtual machine host licensing

Before you start this process, make sure that you have imported the license file and confirmed that it is adequate for managing the hosts that you want to use. The Currently Installed Licenses section includes the total number of CPU licenses, the number being used currently, and the number of licenses remaining. These numbers increment and decrement based on the selections you make on this dialog.

ⓘ | **NOTE:** The Virtual Machine Hosts license applies equally to both Hyper-V and VMware hosts. They are treated the same in terms of licensing.

To manage virtual machine host licensing

1 Click **Tools**, and then click **Options**.

2 Under the **Licensing** node, click **Virtual Machine Hosts**.

3 In the Currently Installed Licenses section, view the **CPUs Allowed** value to confirm the number of sockets that are licensed for protection. The number of licenses currently allocated is shown in the **CPUs Used** column.

4 The Licensed Virtual Machine Hosts section lists the managed hosts, the number of CPUs for each host, and the licensed features for that host.

    a The **CPUs** column indicates the license "cost" of each host.

    b In the **Pro Backup** and/or **Replication** column, select or deselect hosts until you have used all of the available licenses, or until all of your hosts are protected.

5 Click **OK**.

Managing physical machine licensing

Before you start this process, make sure that you have imported the license file and confirmed that it is adequate for managing the number of physical machines you want to protect. The Currently Installed Licenses section includes the total number of physical machine licenses available, and the number being used currently. These numbers increment and decrement based on the selections you make on this dialog.

To manage physical machine licensing

1 Click **Tools**, and then click **Options**.

2 Under the **Licensing** node, click **Physical Machine**.

3 In the Currently Installed Licenses section, view the **Machines Allowed** value to confirm the total number of physical machines licenses available The number of licenses currently allocated is shown in the **Machines Used** column.

4   The Licensed Physical Machine section lists the physical machines in vRanger's inventory, and indicates whether each machine is licensed for backup.

    a   In the **Physical Backup** column, select or deselect machines until you have used all of the available licenses, or until all of your machines are protected.

5   Click **OK**.

Managing job configuration settings

Through the Configuration Options dialog, you can set limits for the number of tasks running, the time frame for task execution, and the space allowed on a datastore to accommodate snapshot growth.

The maximum number of tasks are defined in the following table.

Table 1. Maximum number of tasks

On vRanger	The maximum number of backup, restore and replication tasks running on this instance of vRanger at any one time
off a LUN	The maximum number of tasks allowed on any particular logical storage device
on a host	The maximum number of tasks running on a host at any given time
per repository	The maximum number of tasks allowed on any specific repository at any specific time
locally	The number of backup and restore tasks that are physically executing on the local vRanger machine.
per VA	The number of replication tasks running on a vRanger virtual appliance.

To manage job configuration settings

1   Click **Tools**, and then click **Options**.

2   Under the My Jobs node, click **Configuration**.

3   In the **Resource Management** section, set a maximum for the number of tasks that can be run:

- On vRanger (>=1) – Default is 100.
- Off a LUN (1-5) – Default is 3.
- On a host (1-4) – Default is 1.
- Per repository (>=1) – Default is 3.
- Locally (1 - 20) – Default is 3.
- Per VA (1 -18) – Default is 2.

4   In the **Timeout** section, populate the **Task Timeout** (1-999 hours) text box; Default is 24

5   In the **Space needed on host** (in GBs) text box, enter a number; Default is 6.

6   Click **OK**.

Enabling cataloging options

The vRanger cataloging function must be enabled before data collection can begin. The Catalog dialog includes the option to enable or disable catalog collection, and displays the status of the Catalog Service.

ⓘ | **NOTE:** For Catalog collections to function, the Catalog Service must display a status of "Running."

To enable cataloging options

1   Click **Tools**, and then click **Options**.

2   Under the My Jobs node, click **Catalog**.

3   Select **Enable Catalog Data Collection**.

4   Click **OK**.

> ⓘ | **NOTE:** Cataloging must also be enabled for each backup job for which you want to collect catalog data.

About catalog filtering

While there are thousands (or hundreds of thousands) of files in a typical VM, most are not relevant to file level recovery operations. To streamline cataloging operations, and reduce impact to the catalog database, vRanger filters files to be indexed in two ways:

- **Path** – by default, vRanger does not catalog any files in the directories listed below. Path filtering is determined by entries in the PathFilterTokens.txt file, located at **C:\Program Files\Quest Software\vRanger\Service\Configuration**.

Program Files	$Txf
Windows	RECYCLER
$Extend	System Volume Information
$TxfLog	I386

- **File** – By default, vRanger does not catalog files of the type below. File filtering is determined by entries in the FilesFilterTokens.txt file, located at **C:\Program Files\Quest Software\CatalogManager\Config\Files**.

.lnk	$Boot	$Reparse
$MFT	$BadClus	$RmMetadata
$Volume	$Secure	$Repair
$AttrDef	$UpCase	$Tops
$BitMap	$Quota	$TxfLog
	$ObjID	

> ⓘ | **NOTE:** File filtering applies to un-filtered paths. If a path is filtered, files in that path do not need to be.

For most situations, the default filtering options are sufficient. If you want to filter out additional paths or files, simply add the path or file to the appropriate text file.

Managing Transport Failover Options

Replication transport actions can occur either in the Service Console or through HotAdd (Lan-Free or Network). If one transport method is not available, vRanger attempts to perform the replication task using one of the other transport methods. The order in which you should sort the transport options depends on what type of host you have. The VA-based Network option is the slowest of the three, and should be configured last.

- ESX – The recommended transport failover order for ESX hosts is **VA-based Hot-Add> ESX Console> VA-based Network**.

- ESXi/ESX 5 – For hosts without a service console, configure the transport order to **VA-based HotAdd> VA-based Network**.

To manage transport failover settings

1   Click **Tools**, and then click **Options**.

2 Under the My Jobs node, click **Transport Failover.**

3 Select the transport options available in your environment.

4 To change the order of the transport options, click an option to highlight it and use the up or down arrows to move the item higher or lower in the order.

Supplemental instructions – additional repository types

The main procedures describe the process for adding a CIFS repository. To add an NFS, FTP, SFTP, NetVault SmartDisk, EMC Data Domain Boost (DDB), or Dell Rapid Data Access (RDA) repository, see the appropriate section below:

- Adding an NFS repository
- Adding an FTP repository
- Adding an SFTP repository
- Adding a NetVault SmartDisk repository
- Adding an EMC Data Domain Boost (DDB) repository
- Adding a Dell Rapid Data Access (RDA) repository

ⓘ | **NOTE:** Dell RDA and EMC DDB repositories support repository replication. For more information on configuring repository replication, refer to Managing Repository Replication.

Adding an NFS repository

The section below provides the information required to configure an NFS repository in vRanger.

NFS requirements

The following configurations are required for adding an NFS repository:

- vRanger supports the NFS3 protocol, which does not provide for authentication. vRanger uses the NFSNobody account to connect to the repository.
- The NFSNobody account must have R,W, X permissions for the target directory.
- Your */etc/exports* file must contain the export directory and a list of the IPs that can connect to the repository. You may list each individually (the vRanger machine and any ESX hosts using direct-to-target) or use a wildcard *. The export directory should be shared with read/write asynchronous access. For an export directory of */backups*, this would look like:

```
[root@NFSServer etc]# cat exports
/backups *{rw, async}
```

Adding an NFS repository

Complete the steps in the following procedure to add an NFS repository.

To add an NFS repository

1 To add an NFS repository, navigate to the My Repositories pane.

2 Right-click anywhere in the My Repositories pane, select **Add**, and then click **NFS**.

3 On the **Add Network File Share Repository** dialog, populate the text boxes described below:

- **Repository Name** – A descriptive name for the repository.

- **Description** – A long form description for the repository. This is an optional text box.

- **DNS Name or IP** – The IP or FQDN for the repository.

- **Export Directory** – Specify the Export directory, which is similar in concept to a network share. You need to create a target sub-directory in the export directory.

- **Target Directory** – This is a sub-directory of the NFS Export directory. This is the location to which savepoints are written.

4   Select **Encrypt all backups to this repository** if you want these backups to be password-protected.

5   Enter a Password for the encrypted repository — confirm the password by re-entering it.

⚠ | **WARNING:** **If you want to use the Encrypt all backups feature, make certain to retain the password you enter in the following steps. There is no back-door or admin-level password. If the password is unknown, the backups are not usable.**

6   Click **OK**.

The connection to the repository is tested and the repository is added to the My Repositories pane and the Repository Information dialog.

7   vRanger checks the configured repository location for existing manifest data to identify existing savepoints. If any are found, you are prompted to take one of three actions:

- **Import as Read-Only** – With this option, all savepoint data is imported into the vRanger database, but only for restores. You cannot back up to this repository.

- **Import** – All savepoint data is imported into the vRanger database. vRanger can use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – The savepoint data is retained on the disk, but cannot be imported into vRanger. vRanger ignores the existing savepoint data and treats the repository as new.

Adding an FTP repository

The process for adding an FTP repository is described below.

To add an FTP repository

1   To add an FTP repository, navigate to the My Repositories pane.

2   Right-click anywhere in the My Repositories pane, select **Add**, and then click **FTP**.

3   On the **Add FTP Repository** dialog, populate the text boxes described below:

- **Repository Name** – a descriptive name for the repository.

- **Description** – a long form description for the repository. This is an optional text box.

- **DNS Name or IP** – the IP or FQDN for the repository

- **Target Directory** – this is the directory on the FTP server to which the repository is written.

- **User name** – the user name for the account that vRanger uses to connect to the repository.

- **Password** – the password for the user account above.

- **Port** – the port that vRanger uses to connect to the FTP server. The default FTP port is 21

- **Timeout** – this value determines how long to wait before terminating an idle connection. The default value is 600 seconds

4   Select **Encrypt all backups to this repository** if you want these backups to be password-protected.

5   Enter a Password for the encrypted repository — confirm the password by re-entering it.

> ⚠ **WARNING:** **If you want to use the Encrypt all backups feature, make certain to retain the password you enter in the following steps. There is no back-door or admin-level password. If the password is unknown, the backups are not usable.**

6   Click **OK**.

The connection to the repository is tested and the repository is added to the My Repositories pane and the Repository Information dialog.

7   vRanger checks the configured repository location for existing manifest data to identify existing savepoints. If any are found, you are prompted to take one of three actions:

-    **Import as Read-Only** – With this option, all savepoint data is imported into the vRanger database, but only for restores. You cannot back up to this repository.

-    **Import** – All savepoint data is imported into the vRanger database. vRanger can use the repository for backups and restores. vRanger requires read and write access to the directory.

-    **Overwrite** – The savepoint data is retained on the disk, but cannot be imported into vRanger. vRanger ignores the existing savepoint data and treats the repository as new.

Adding an SFTP repository

The process for adding an SFTP repository is described below.

To add an SFTP repository

1   To add an SFTP repository, navigate to the My Repositories pane.

2   Right-click anywhere in the My Repositories pane, select **Add**, and then click **SFTP**.

3   On the **Add SFTP Repository** dialog, populate the text boxes described below:

-    **Repository Name** – A descriptive name for the repository.

-    **Description** – A long form description for the repository. This is an optional text box.

-    **DNS Name or IP** – The IP or FQDN for the repository.

-    **Target Directory** – This is the directory on the SFTP server to which the repository is written.

-    **User name** – The user name for the account that vRanger uses to connect to the repository.

-    **Password** – The password for the user account above.

-    **Port** – The port that vRanger uses to connect to the SFTP server. The default SFTP port is 22.

-    **Timeout** – This value determines how long to wait before terminating an idle connection.

4   Select **Encrypt all backups to this repository** if you want these backups to be password-protected.

5   Enter a Password for the encrypted repository — confirm the password by re-entering it.

> ⚠ **WARNING:** **If you want to use the Encrypt all backups feature, make certain to retain the password you enter in the following steps. There is no back-door or admin-level password. If the password is unknown, the backups are not usable.**

6   Click **OK**.

The connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

7   vRanger checks the configured repository location for existing manifest data to identify existing savepoints. If any are found, you are prompted to take one of three actions:

-    **Import as Read-Only** – With this option, all savepoint data is imported into the vRanger database, but only for restores. You cannot back up to this repository.

- **Import** – All savepoint data is imported into the vRanger database. vRanger can use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – The savepoint data is retained on the disk, but cannot be imported into vRanger. vRanger ignores the existing savepoint data and treats the repository as new.

Adding a NetVault SmartDisk repository

The process for adding a NetVault SmartDisk repository is described below.

To add a NetVault SmartDisk repository

1   To add a NetVault SmartDisk repository, navigate to the My Repositories pane. Right-click anywhere in the My Repositories pane. Select **Add--> NVSD**.

2   On the **Add NVSD Repository** dialog, populate the fields below:

    a   **Repository Name** – a descriptive name for the repository.

    b   **Description** – a long form description for the repository. This is an optional field.

    c   **DNS Name or IP** – the IP or FQDN for the NetVault SmartDisk server.

    d   **User name/Password** – the credentials for the account that vRanger will use to connect to the repository.

> ⓘ **NOTE:** If you have enabled WebDav authentication on your NetVault SmartDisk server, then use those credentials here. If you have not enabled WebDav authentication, then no credentials are required. See the *Dell NetVault SmartDisk Administrator's Guide* for more information.

    e   **Target Directory** – this value is a logical identifier for the repository on the NetVault SmartDisk Instance, rather than a specific directory on the NetVault SmartDisk server. If you have previously created one or more repositories on the target NetVault SmartDisk Instance, the existing target directories will be available in a drop-down menu.

        To create a new repository, enter a name for your repository.

    f   **Port** – the port that vRanger will use to connect to the NetVault SmartDisk server. The default port is 37453.

    g   **Timeout** – this value determines how long to wait before terminating an idle connection.

> ⓘ **NOTE:** Note that encryption is not supported for NetVault SmartDisk repositories.

3   Click **Ok**.

4   Click **OK** -- the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

5   vRanger will check the configured repository location for existing manifest data to identify existing savepoints. If any are found, you will be prompted to take one of three actions:

- **Import as Read-Only** – with this option, all savepoint data will be imported into the vRanger database, but only for restores. You will not be able to back up to this repository.

- **Import** – all savepoint data will be imported into the vRanger database. vRanger will be able to use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – the savepoint data will be retained on the disk, but will not be imported into vRanger. vRanger will ignore the existence of the existing savepoint data and treat the repository as new.

Removing a NetVault SmartDisk repository

The process for removing a NetVault SmartDisk repository is the same as removing any other repository type.

To remove a NetVault SmartDisk repository

1 From the My Repositories view, right-click on the NetVault SmartDisk repository and select **Remove**.

2 The **Remove Repository** dialog appears, showing the savepoints in the selected repository.

   When removing a repository, you have the option of keeping the savepoints on disk or deleting them.

3 To remove the savepoints, select **Delete all savepoints in this repository**.

   OR

   To keep the savepoints on disk, ensure that **Delete all savepoints in this repository is not selected**.

⚠ | **WARNING: This will delete the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.**

4 If you chose to delete the savepoints, the **Delete Savepoints** dialog appears. Click **OK**. The savepoints will be deleted, along with the repository.

Adding an EMC Data Domain Boost (DDB) repository

The process for adding an EMC Data Domain Boost repository is described below. This procedure assumes the following:

- That you have a properly configured Data Domain appliance that is accessible to the vRanger machine.

- That DD Boost is licensed and enabled on the Data Domain device.

- That you have designated a DD Boost User account.

ⓘ | **NOTE:** Consult the Data Domain documentation for information on performing the configurations listed above.

To add an EMC Data Domain Boost Repository

1 To add a DD Boost repository, navigate to the My Repositories pane. Right-click anywhere in the My Repositories pane. Select **Add--> EMC Data Domain Boost (DDB)**.

2 On the **Add EMC Data Domain Boost Repository** dialog, populate the fields below:

   a **Repository Name** – a descriptive name for the repository.

   b **Description** – a long form description for the repository. This is an optional field.

   c **DNS Name or IP** – the IP or FQDN for the Data Domain device.

   d **User name/Password** – these are the credentials for the account selected as the **DD Boost User** on the Data Domain device.

   ⓘ | **NOTE:** Consult the Data Domain documentation for more information on setting a DD Boost user name.

3 **Storage Unit** – this is the Storage Unit configured on the Data Domain device. Select one from the drop-down menu, or enter the name for a new Storage Unit and vRanger will create one for you. Click **OK** -- the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

   ⓘ | **IMPORTANT:** The supported characters are alphanumeric characters and ~!@#$^_+`-={}[],.

4 vRanger will check the configured repository location for existing manifest data to identify existing savepoints. If any are found, you will be prompted to take one of three actions:

- **Import as Read-Only** – with this option, all savepoint data will be imported into the vRanger database, but only for restores. You will not be able to back up to this repository.

- **Import** – all savepoint data will be imported into the vRanger database. vRanger will be able to use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – the savepoint data will be retained on the disk, but will not be imported into vRanger. vRanger will ignore the existence of the existing savepoint data and treat the repository as new.

Configuring Repository Replication

Managed File Replication is configured through the My Repositories view of the vRanger interface. The procedure below assumes that the following are true:

- You have configured at least one DD Boost repository in vRanger.

- At least one other Data Domain device with DD Boost enabled is available in your environment to serve as the replication target. This device does not need to be added to vRanger.

- All devices used in replication have a Replicator license installed.

To configure Managed File Replication

1 From the My Repositories view, right-click on the Data Domain Boost repository and select **Configure Replication**.

2 On the **Configure Replication** dialog, populate the fields below:

   a **Repository Name** – a descriptive name for the repository.

   b **Description** – a long form description for the repository. This is an optional field.

   c **DNS Name or IP** – the IP or FQDN for the Data Domain device.

   d **DD Boost User name/Password** – these are the credentials for the account selected as the **DD Boost User** on the Data Domain device.

   e **Storage Unit** – this is the Storage Unit configured on the Data Domain device. Select one from the drop-down menu, or enter the name for a new Storage Unit and vRanger will create one for you.

3 Click **OK** -- the connection to the device is tested and the device is added as a repository is added to the My Repositories pane and the Repository Information dialog.

Once Managed File Replication is configured for a repository, the **Configure Replication** option will be disabled for that repository.

ⓘ | IMPORTANT: Once a repository is configured for replication, you must select a synchronization method before replication will occur. Refer to Managing Repository Replication for information on manual and automatic synchronization options.

Editing a replication configuration

You may edit an existing replication configuration to update credentials or timeout values. You may also use the **Edit Repository Details** dialog to view free space for the repository.

1 From the My Repositories View, right-click on the **EMC Data Domain Boost** repository and select **Edit Replication Configuration**.

2 On the **Edit EMC Data Domain Boost Repository** dialog, you may edit the following fields:

- Repository Name

- User Name

- Password

3 Alternatively, view the **Free Space** field for up-to-date information about this repository.

4   Make the desired changes, and click Update.

ⓘ | **NOTE:** For information on disabling repository replication, refer to Disabling repository replication.

Removing an EMC Data Domain Boost (DDB) repository

The process for removing a Data Domain Boost repository is the same as removing any other repository type, except that you also have the option of removing the storage unit used by the repository. There is a hard limit on the number storage units that can be configured on a Data Domain device, which makes removing unused storage units an important maintenance effort.

To remove a DD Boost repository

1   From the My Repositories view, right-click on the Data Domain Boost repository and select **Remove**.

2   The **Remove Repository** dialog appears, showing the savepoints in the selected repository.

     When removing a repository, you have the option of keeping the savepoints on disk or deleting them. In order to remove the storage unit associated with the repository, you will need to remove the savepoints.

     If Managed File Replication is configured for this repository, you will also be given the option to delete the savepoints in the replicated repository as well.

3   Select **Delete all savepoints in this repository**. To delete the savepoints in the replicated repository, select **Delete all savepoints in replication repository**. Click **OK**.

⚠ | **WARNING: This will delete the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.**

4   The **Delete Savepoints** dialog appears. Select **Delete DD Boost storage unit**, and click **OK**. The savepoints will be deleted, along wit the storage unit associated with the repository.

Adding a Dell Rapid Data Access (RDA) repository

The Dell DR Series disk-based data protection appliances optimize storage utilization and reduce network bandwidth requirements with in-line deduplication, server -side compression, and compressed and deduplicated replication.

The DR Series include the following appliance types:

- **DR4000 System** – This consists of preinstalled DR System software on a Dell R510 appliance platform.
- **DR4100 System** – This consists of preinstalled DR System software on a Dell R720xd appliance platform.
- **DR6000 System** - This is preinstalled DR6000 system software on a Dell PowerEdge R720xd appliance platform. This differs from the DR4100 by including a higher level of base system hardware.

The Dell DR Series supports CIFS, NFS, and RDA protocols. The RDA protocol provides a logical disk interface for the DR Series system, and enables better coordination and integration between vRanger and the DR Series system and provides for client-side deduplication of vRanger backups.

ⓘ | **NOTE:** It is recommended to use the RDA protocol when using a Dell DR Series system as a repository.

For more information about the Dell DR Series systems, refer to the *Dell DR Series System Administration Guide*.

Adding a Dell DR Series system as an RDA repository

The process for adding a Dell Rapid Data Access (RDA) Repository is described below. This procedure assumes the following:

- That you have a properly configured the Dell DR Series appliance that is accessible to the vRanger machine.

- That you have created at least one storage container to be used as a Logical Storage Unit. When creating the storage container, use the options:

    - Connection type – RDA

    - RDA type – RDS

- That you have designated a RDA User account.

- That you have configured your firewall to enable the following TCP ports.

    - 9904

    - 9911

    - 9915

    - 9916

    - 9920

ⓘ **NOTE:** Consult the Dell DR Series documentation for information on performing the configurations listed above.

To add a Dell DR Series system as a Dell RDA Repository

1 Navigate to the My Repositories pane. Right-click anywhere in the My Repositories pane. Select **Add-->Dell Rapid Data Access (RDA)**.

2 In the Add Dell Rapid Access Repository dialog, configure the following parameters:

    - **Repository Name** – a descriptive name for the repository.

    - **Description** – a long form description for the repository. This is an optional field.

    - **DNS Name or IP** – the DNS name or IP address of the Dell DR Series system.

    - **RDA Username** – a user account that can be used to log on to the device. On the Dell DR Series system, only one user account exists, and the user ID for that account is **"backup_user"**. You can only change the password for this account; you cannot create a new account or delete the existing account.

    - **RDA Password** – the password for the user account.The default password is "St0r@ge!"

    - **Logical Storage Unit** – Enter the name of the storage container. Ensure that the container is created before you add the device. The device will not be added if the specified container does not exist on the device. When creating the storage container, use the options:

        - Connection type – RDA

        - RDA type – RDS

    - **Port Number** – leave this value at 0 to use the default port.

3 vRanger will check the configured repository location for existing manifest data to identify existing savepoints. If any are found, you will be prompted to take one of three actions:

    - **Import as Read-Only** – with this option, all savepoint data will be imported into the vRanger database, but only for restores. You will not be able to back up to this repository.

    - **Import** – all savepoint data will be imported into the vRanger database. vRanger will be able to use the repository for backups and restores. vRanger requires read and write access to the directory.

    - **Overwrite** – the savepoint data will be retained on the disk, but will not be imported into vRanger. vRanger will ignore the existence of the existing savepoint data and treat the repository as new

Configuring repository replication

Data Replication is configured through the My Repositories view of the vRanger interface.

> ⓘ | **TIP:** For more information on repository replication, refer to the *Dell DR Series System Administration Guide*.

The procedure below assumes that the following are true:

- You have configured at least one Dell RDA repository in vRanger.

- At least one other Dell DR Series appliance with RDA is available in your environment to serve as the replication target. This device does not need to be added to vRanger.

To configure Replication

1   From the My Repositories view, right-click on the Dell Rapid Data Access (RDA) repository and select **Configure Replication**.

2   On the **Configure Replication** dialog, populate the fields below:

- **Repository Name** – a descriptive name for the repository.

- **Description** – a long form description for the repository. This is an optional field.

- **DNS Name or IP** – the DNS name or IP address of the Dell DR Series system.

- **RDA Username** – a user account that can be used to log on to the device. On the Dell DR Series system, only one user account exists, and the user ID for that account is **"backup_user"**. You can only change the password for this account; you cannot create a new account or delete the existing account.

- **RDA Password** – the password for the user account.

- **Logical Storage Unit** – Enter the name of the storage container. Ensure that the container is created before you add the device. The device will not be added if the specified container does not exist on the device.When creating the storage container, use the options:

  - Connection type – RDA

  - RDA type – RDS

- **Port Number** – leave this value at 0 to use the default port.

3   Click **OK** -- the connection to the device is tested and the device is added as a repository is added to the My Repositories pane and the Repository Information dialog.

Once replication is configured for a repository, the **Configure Replication** option will be disabled for that repository.

> ⓘ | **IMPORTANT:** Once a repository is configured for replication, you must select a synchronization method before replication will occur. Refer to Managing Repository Replication for information on manual and automatic synchronization options.

Editing a replication configuration

You may edit an existing replication configuration to update credentials or timeout values. You may also use the **Edit Repository Details** dialog to view free space for the repository.

1   From the My Repositories View, right-click on the Dell Rapid Data Access (RDA) repository and select **Edit Replication Configuration**.

2   On the **Edit Dell Rapid Data Access (RDA)** dialog, you may edit the following fields:

- Repository Name

- User Name

- Password

3   Alternatively, view the **Free Space** field for up-to-date information about this repository.

4 Make the desired changes, and click **Update**.

ⓘ | **NOTE:** For information on disabling repository replication, refer to Disabling repository replication.

Removing a Dell Rapid Data Access (RDA) repository

The process for removing a Dell Rapid Data Access (RDA) repository is the same as removing any other repository type.

To remove a Dell Rapid Data Access (RDA) repository

1 From the My Repositories view, right-click on the Dell Rapid Data Access (RDA) repository and select **Remove**.

2 The **Remove Repository** dialog appears, showing the savepoints in the selected repository.

When removing a repository, you have the option of keeping the savepoints on disk or deleting them. In order to remove the storage unit associated with the repository, you will need to remove the savepoints.

If replication is configured for this repository, you will also be given the option to delete the savepoints in the replicated repository as well.

3 Select **Delete all savepoints in this repository**. To delete the savepoints in the replicated repository, select **Delete all savepoints in replication repository**. Click **OK**.

⚠ | **WARNING: This will delete the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.**

4 The **Delete Savepoints** dialog appears. Click **OK**. The savepoints will be deleted.

Adding a Dell DR Series System as a CIFS repository

ⓘ | **NOTE:** Adding a Dell DR Series system as a CIFS repository does not take advantage of the full function of the DR Series system. Dell Software recommends adding the DR Series as an RDA Repository.

To add a Dell DR Series System as a CIFS repository

1 Navigate to the My Repositories pane. Right-click anywhere in the My Repositories pane. Select **Add--> CIFS**.

2 In the **Repository Type** section, select **Windows Share (CIFS)**, if necessary.

3 Populate the **Repository Name** field.

This value will display in the **My Repositories** pane.

4 Populate the **Description** field.

5 Enter a **username** and **password** in the appropriate fields.

6 In the **Server** field, type the UNC path to the desired repository directory. Alternatively, you may enter a partial path and click **Browse** to find the target directory.

ⓘ | **NOTE:** You must enter a valid username and password before using the browse functionality.

7 Do not select **Encrypt all backups to this repository**.

ⓘ | **NOTE:** Using encryption or compression with deduplicated repositories will limit or disable deduplication. Encryption and compression should not be used with any repository type that provides deduplication.

8 Click **Save** -- the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

9 vRanger will check the configured repository location for existing manifest data to identify existing savepoints. If any are found, you will be prompted to take one of three actions:

- **Import as Read-Only** – with this option, all savepoint data will be imported into the vRanger database, but only for restores. You will not be able to back up to this repository.

- **Import** – all savepoint data will be imported into the vRanger database. vRanger will be able to use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite** – the savepoint data will be retained on the disk, but will not be imported into vRanger. vRanger will ignore the existence of the existing savepoint data and treat the repository as new.

10 Click **Next**.

Adding a Dell DR Series System as a NFS Repository

ⓘ **NOTE:** Adding a Dell DR Series system as a NFS repository does not take advantage of the full function of the DR Series system. Dell Software recommends adding the DR Series as an RDA Repository.

To add a Dell DR Series System as an NFS repository

1 Navigate to the My Repositories pane. Right-click anywhere in the My Repositories pane. Select **Add--> NFS**.

2 On the **Add Network File Share Repository** dialog, populate the fields below:

   a **Repository Name** – a descriptive name for the repository.

   b **Description** – a long form description for the repository. This is an optional field.

   c **DNS Name or IP** – the IP or FQDN for the repository

   d **Export Directory** – specify the Export directory, which is similar in concept to a network share. You will need to create a target sub-directory in the export directory.

   e **Target Directory** – this is a sub-directory of the NFS Export directory. This is the location to which savepoints will be written.

3 Do not select **Encrypt all backups to this repository.**

   ⓘ **NOTE:** Using encryption or compression with deduplicated repositories will limit or disable deduplication. Encryption and compression should not be used with any repository type that provides deduplication.

4 Click **OK** -- the connection to the repository is tested and the repository is added to the My Repositories pane and the **Repository Information** dialog.

5 vRanger will check the configured repository location for existing manifest data to identify existing savepoints. If any are found, you will be prompted to take one of three actions:

- **Import as Read-Only** – with this option, all savepoint data will be imported into the vRanger database, but only for restores. You will not be able to back up to this repository.

- **Import** – all savepoint data will be imported into the vRanger database. vRanger will be able to use the repository for backups and restores. vRanger requires read and write access to the directory.

6 **Overwrite** – the savepoint data will be retained on the disk, but will not be imported into vRanger. vRanger will ignore the existence of the existing savepoint data and treat the repository as new.

Understanding the vRanger virtual appliance

The vRanger virtual appliance is a small, pre-packaged Linux distribution that servers as a platform for vRanger operations away from the vRanger server. vRanger uses the virtual appliance for the following VMware functions:

- Replication to and from ESXi hosts

- File-level recovery from Linux machines

- Optionally, for backups and restores

You may also use the VA on ESX hosts as well to avoid taxing the Service Console resources, which are limited.

> ⓘ | **NOTE:** The virtual appliance is not designed to be used with Hyper-V machines.

For more information about vRanger virtual appliance configuration, see the following topics:

- Virtual appliance overview

- The virtual appliance scratch disk

- Strategies for sizing the scratch disk

- Deploying the virtual appliance by using the Virtual Appliance Deployment Wizard

- About deploying multiple virtual appliances

- Changing the virtual appliance configuration

- Process of configuring virtual appliance networking

- Disabling DRS for the vRanger virtual appliance

Virtual appliance overview

The vRanger virtual appliance is a lightweight Linux virtual machine that serves as an operating environment for the components that would normally run in the Service Console. Removing the workload from the Service Console is clearly necessary for ESXi and vSphere 5, but it is advantageous for ESX servers as well. While the Service Console is restricted to a single CPU (CPU 0), with virtual appliances the workload can be spread across the other CPUs available to a host. This provides increased reliability and scalability over operations limited to the Service Console.

The virtual appliance must be deployed to any ESXi host that you want to configure for replication (either as a source or a destination). For hosts in a cluster, you may deploy just one VA to the cluster; the VA is shared among the cluster's hosts. When deploying a VA to a cluster, you must chose a host in the cluster to which the VA should be associated.

In addition, replication by way of a virtual appliance requires that if a VA is used on one host or cluster in a replication job, then a VA must be used on both the source and destination host or cluster. In other words, virtual appliances, when used for replication, must be used in pairs.

When configuring the virtual appliance, special consideration should be given to the amount of resources (CPU and RAM) allocated to the VA as the number of simultaneous tasks the VA can process is directly tied to available resources. In addition, if you want to perform replication tasks using a virtual appliance, carefully consider an appropriate size for the VA scratch disk. See the section below for more information on scratch disk usage.

The virtual appliance scratch disk

If you want to perform replication tasks with your virtual appliance, you need to add a second "scratch" disk to the VA. This scratch disk is used to store two types of files:

- *vzmap* **files** – block maps (in the form of a *vzmap* file) for the VMs replicated to the destination host. This is block map information, and not actual data blocks. These maps are compared to the source VM

during each replication to identify the data blocks that have changed since the last replication. The *vzmap* files make differential replication significantly faster as they remove the need to scan the destination VM blocks for comparison with the source VM.

- **vzUndo files** – As data is sent to the destination host (via the VA), blocks in the destination disk are written to the undo file before they are overwritten by the changed data. If replication fails and an undo becomes necessary, then the original destination disk blocks are read from the undo file and written to the destination disk to rollback the failed replication. This is a key function designed to provide resiliency in the face of a network failure; if there is a network failure during the replication pass, the destination VM is not corrupted by incomplete data.

After the replication is complete, and all data has been received by the destination VA, the undo file is deleted. At that point, the storage space used by the undo file is returned to the VA for use. Note that undo files are not created during the first replication. During the first replication, the entire VM is sent to the destination host, but there is no existing data on the destination VMDKs, and therefore no risk of corruption. Data is streamed directly to the VMDK. You do not need to allocate scratch disk space for this scenario.

While the vzmap files are trivial in size (in the order of a few MBs), the undo file can potentially be as large as the VM itself. While the scratch disk needs to be configured to a size sufficient to handle the data of concurrent replication tasks, making it too large wastes valuable storage space. Use the sections below to guide you in determining the proper size for the scratch disk.

Strategies for sizing the scratch disk

The scratch disk needs to be large enough only to hold the permanent *vzmap* files and the temporary *vzUndo* files, plus a small margin for safety. How large that is depends almost entirely on the amount of changed data you are replicating. The amount of changed data is itself a function of the number of VMs you'll be replicating, their total disk size, replication frequency, and the data change rate per VM. It is important to understand all of this data when sizing the scratch disk.

If you are using one VA for a cluster, remember that you must account for all simultaneous replications for the cluster.

Use historical data

If you have previously replicated the source VMs, the most accurate method to properly size the scratch disk (without wasting storage space) is to use historical replication data available in the Replicate Task Reports (in the vRanger My Reports view) for the VMs in question. This report shows the amount of data written during each replication task.

The safest method to size your scratch disk based on historical data is to record the highest amount of data written for each VM that you replicate at one time, and size the disk to accommodate those values.

To avoid filling your scratch disk, Dell recommends that you add a small margin (10% or so) to the calculated scratch disk size for safety.

Calculating

If you do not have information on the amount of changed data for each VM, you can estimate the appropriate size of the scratch disk based on the VM size and the number of VMs you plan to replicate at one time.

A general rule for sizing the scratch disk is to choose a percentage of the total VM size to represent the practical limit of changed data for a given replication. Only you can decide what is appropriate for your environment. The numbers below are examples given to illustrate the concept:

For example, if you have four virtual machines that you want to replicate to a host or cluster at the same time, the minimum requirements for the VMs are described in the following table.

Table 2. Minimum requirements

VM	VM Size	Change Rate	Change Size
1	100 GB	15%	15 GB
2	100 GB	10%	10 GB

Table 2. Minimum requirements

VM	VM Size	Change Rate	Change Size
3	100 GB	20%	20 GB
4	60 GB	5%	3 GB

For the VMs above, we would need approximately 48 GB of disk space for the undo files, plus a buffer of approximately 10% (for safety's sake). In our example, an appropriate estimate for the scratch disk size for the VMs above would be approximately 55 GB.

Bear in mind that the estimate exercise should be done for every set of VMs you want replicated to that host or cluster, with the scratch disk being sized to accommodate the largest value obtained.

⚠ | **WARNING: If the scratch disk runs out of space, replication tasks fail with the error "FATAL cant_write can't write vix".**

Options for a smaller scratch disk

As previously stated, the primary driver for the scratch disk size is the amount of changed data that needs to be replicated. If you need to reduce the storage requirements for your scratch disk, you can:

- Reduce the amount of time that passes between replication intervals. More frequent replications contain smaller amounts of changed data.

- Reduce the number of VMs that you replicate simultaneously. Doing so sends less data through the scratch disk at any one time, which requires less space. Remember that the *vzundo* files are deleted after the replication completes.

- When provisioning the scratch disk, select the Thin Provisioning option. Thin provisioning helps prevent wasted storage space by not allocating blocks until they are written by the VM.

The scratch disk on the source host

As the scratch disk is used primarily for staging changes before they are written to disk – activity which occurs on the destination host or cluster– the scratch disk on the source side can be kept fairly small. However, in the event that you need to fail over to the D/R site, the replication job reverses direction and starts replicating changes back to the product site (the original source host or cluster). For this to occur, the scratch disk on the source side needs to be re-sized to accommodate the changed data.

Scratch disk location

When creating the second disk, make sure you place the disk on a datastore with block sizes large enough to support the expected VMDK. The list below shows the maximum file size available for each block size:

- 1 MB block size – 256 GB maximum file size

- 2 MB block size – 512 GB maximum file size

- 4 MB block size – 1024 GB maximum file size

- 8 MB block size – 2048 GB maximum file size

ⓘ | **NOTE: This limitation does not apply to VMFS 5.**

Deploying the virtual appliance by using the Virtual Appliance Deployment Wizard

vRanger uses a virtual appliance for replication to and from ESXi hosts, for file-level recovery from Linux machines, and optionally for backups and restores.The VA Deployment Wizard offers a simple method to deploy

virtual appliances one at a time. You may, of course, run the wizard for each VA that you need to deploy, or alternatively deploy multiple VAs at once as described in About deploying multiple virtual appliances.

> (i) **NOTE:** The automated deployment process requires DHCP. To deploy a virtual appliance through the VA Deployment Wizard, you must have a DHCP server accessible to the virtual appliance during deployment.

To deploy the virtual appliance (VA) using the Virtual Appliance Deployment Wizard, complete the following tasks:

1 Launching the Virtual Appliance Deployment Wizard

2 Selecting the Virtual Appliance OVA file

3 Selecting a VA host

4 Selecting VA deployment options

5 Configuring the VA IP address

6 Confirming VA deployment selections

Launching the Virtual Appliance Deployment Wizard

Complete the steps in the following procedure to launch the Virtual Appliance Deployment Wizard.

To launch the Virtual Appliance Deployment Wizard

1 Do one of the following:

- From the Tools menu, click **Virtual Appliance Deployment Wizard**.

- From the My Inventory view, right-click on the host to which the VA should be deployed, and click **Deploy Virtual Appliance**.

- In the Virtual Appliance Configuration node of the **Configuration Options** dialog, click **Deploy Virtual Appliance**.

The Virtual Appliance Deployment Wizard appears.

2 Click **Next**.

Selecting the Virtual Appliance OVA file

The OVA file for the vRanger virtual appliance is installed with the vRanger application. The Virtual Appliance Deployment Wizard chooses this OVA by default.

To select the Virtual Appliance OVA file

1 Do one of the following:

- To use the default OVA, click **Next**.

- If you need to use a different OVA, click **Browse** and browse to the location to which you have saved the preferred OVA file.

2 Click **Next**.

Selecting a VA host

You may deploy a VA to a single host, or to a cluster. When performing a backup, restore, or replication task, vRanger first checks for a virtual appliance on the host.If no VA is associated with the host, and if the host is part of a cluster, vRanger checks for a VA on the cluster.

To select a VA host

1. Select **Deploy VA on Cluster**. Expand the preferred cluster to select a host in the cluster to which the VA should be associated.

   Select **Deploy VA on specific host**. Select the host to which the VA should be deployed.

2. Click **Next**.

Selecting VA deployment options

In the VA Deployment Options dialog, you can configure the virtual appliance's name and allocated resources. Additionally, you can configure the size of the VA's scratch disk. See The virtual appliance scratch disk for more information.

To select VA deployment options

1. In the Virtual Appliance Properties section, confirm the VA Name. Edit as required.

2. In the **Virtual Appliance Option** section, configure the resources allocated to the virtual appliance. You may select one of two preconfigured options, or configure a custom setting.

   - **Minimum Required** – This setting allocates one CPU and 512 MB of RAM. This is usually sufficient for two concurrent tasks per VA.

   - **Dell Recommended** – This setting allocates two CPUs and 1 GB of RAM. This is sufficient for four concurrent tasks per VA.

   - **Dell Recommended (with RDA repository)** - This setting allocates four CPUs and 2 GB of RAM. This is sufficient for running tasks going to RDA repositories.

   - **Custom Setting** – Select this value to configure the VA with higher resources for five or more concurrent tasks per VA.

3. In the **VA Datastore** text box, select the datastore for the VA's primary disk.

4. In the **Network Assignment** text box, select the network for the VA's primary NIC.

5. If you want to be able to perform replication tasks on this virtual appliance, you need to configure a scratch disk on the VA. See The virtual appliance scratch disk for more information.

   Select **Use this virtual appliance for replication** and configure the scratch disk size and datastore location. Use the information in Strategies for sizing the scratch disk to guide you.

   > ⓘ **NOTE:** If are upgrading an existing virtual appliance and want to migrate the scratch disk, do not select this option. Refer to Upgrading the vRanger virtual appliance for more information.

6. In the **VA Password** and **Confirm password** text box, enter a new password for the virtual appliance. If you change the password, this becomes the default password for subsequent VA deployments performed during this session.

   OR

   Leave the default password of "vzroot1".

   > ⓘ **NOTE:** You may change the password in the future as described in Changing the Virtual Appliance Configuration.

7. If you want to perform file-level recovery from backups of Linux virtual machines, you need to configure a VA to use for Linux file-level recovery. Select **Use virtual appliance for Linux File Level Restore**.

   If you select this option for a different VA in the future, then that VA is used for Linux file-level recovery.

Configuring the VA IP address

In the VA IP Address Configuration page, you must configure the network configuration for the virtual appliances primary NIC.

To configure the VA IP address

1 Select the method to obtain the IP address:

- **Use DHCP for IP assignment** automatically assigns IP settings to the VA if a DHCP server is available.

- **Use Static IP** lets you deploy the VA with a manual IP configuration.

   ⓘ **NOTE:** If you need to change your network configuration after the VA is deployed, you need to do so from within the virtual appliance. Refer to Process of Configuring Virtual Appliance Networking for instructions.

2 Select your preferred method for DNS resolution.

- **Obtain the DNS server address automatically** uses the DNS settings provided by your DHCP server.

- **Use the following DNS server address** lets you specify DNS settings manually.

3 Click **Next**.

Confirming VA deployment selections

Complete the steps in the following procedure to confirm the VA deployment options displayed are correct.

To confirm VA deployment selections

1 If the configuration options are not correct, click **Back** to make changes to the information.

2 Optionally, select the option to **Power on the VA after deployment is complete**, and then see Creating a template.

3 If the configuration options are correct, click **Finish**.

Creating a template

Optionally, if you need to deploy the vRanger virtual appliance to multiple hosts, it may be helpful to create a virtual machine template from the configured virtual appliance.

To create a template

1 From the VI Client, right-click the configured virtual appliance, select **Template**, and then click **Clone to Template**.

2 Enter the template name and select a location. Click **Next**.

3 Select the host or cluster on which the template is to be stored. Click **Next**.

4 Select the datastore in which the template is to be stored. Click **Next**.

5 On the **Disk Format** dialog, select **Same format as source**. Click **Next**.

6 Review the configurations and click **Finish**.

About deploying multiple virtual appliances

The virtual appliance must be deployed to any ESXi host that you want to configure for replication (either as a source or a destination). For hosts in a cluster, you may deploy the VA to just one host in the cluster; the VA is shared among the cluster's hosts.

In addition, replication by way of a virtual appliance requires that if a VA is used on one host or cluster in a replication job, then a VA must be used on both the source and destination host or cluster. In other words, virtual appliances, when used for replication, must be used in pairs

For a small number of virtual appliances, using the vRanger interface to deploy the VAs is sufficient. To streamline the process of deploying a a high number of virtual appliances, Quest recommends that you deploy and configure the virtual appliance once, then save it as a template to be used for additional deployments using VMware's PowerCLI. The Creating a template procedures include an optional step for creating a template from the configured VA. See the document Deploying Multiple vRanger Virtual Appliances with PowerCLI on the vRanger Community site for more information on using PowerCLI to deploy multiple VAs from a template.

Changing the virtual appliance configuration

After a virtual appliance is deployed, you may change its configuration from the vRanger interface.

> ⓘ **IMPORTANT:** Reconfiguring the virtual appliance may cause it to be power cycled, possibly affecting configured jobs. You are prompted with the jobs configured for the virtual appliance; review these carefully to avoid service interruption.

To change the virtual appliance configuration

1   Do one of the following:

   - From the My Inventory view, right-click the virtual appliance you want to change, and then click **Virtual Appliance Configuration**.

   - From the vRanger toolbar, do the following:

      - Click **Tools**, and then click **Options**.

      - On the Configuration Options dialog, select **Virtual Appliance**.

      - Under Configure Existing Virtual Appliances, select the virtual appliance to change, and then click **Edit**.

   The **Modify Virtual Appliance Configuration** dialog appears.

2   In the Virtual Appliance Properties section, do the following:

   a   Next to **VA Name**, confirm the name of the virtual appliance.

   b   Next to **IP Address**, confirm the IP vRanger uses to connect to the VA.

   By default, vRanger connects to the vRanger virtual appliance using the first IP address reported by the vSphere API. This IP is displayed under Virtual Appliance Properties in the IP Address text box. If you have only one NIC configured on the virtual appliance, continue to the next step.

   If you have more than one NIC configured on the virtual appliance, you may not want vRanger to connect to the first NIC. Select **Override IP Address** to configure vRanger with the IP address for the NIC to which vRanger should connect.

   For example, if you have two NICs configured on the virtual appliance, and want vRanger to connect to the second NIC, select **Override IP Address** to enable the **IP Address** text box, and then enter the IP for the second NIC in the IP Address text box.

   > ⓘ **NOTE:** For information about manual network configuration procedures, see Process of Configuring Virtual Appliance Networking.

3   In the Virtual Appliance Options section, use the **Select an option** drop-down list to change the resources allocated to the virtual appliance as needed. You may select from one of the following options:

   - **Minimum Requirement** – This setting allocates one CPU and 512 MB of RAM. This is usually sufficient for two concurrent tasks per VA.

   - **Quest Recommended** – This setting allocates two CPUs and 1 GB of RAM. This is sufficient for four concurrent tasks per VA.

   - **Custom Setting** – Select this value to configure the VA with your own preferred settings; for example, higher resources for more concurrent tasks per VA.

4   To change the datastore for the VA's primary disk, select the new datastore from the **VA Datastore** drop-down.

> **NOTE:** If you are using a vSphere Virtual Center version below 5.0, the ability to change the datastore of the virtual appliance scratch disk used for replication is not available.

5   If you want to perform replication tasks on this virtual appliance, you need to configure a scratch disk on the VA.

  • Select **Use this virtual appliance for replication** and configure the scratch disk size using the up-and-down arrows and the drop-down list. For more information about sizing a scratch disk, see Strategies for sizing the scratch disk.

For more information, see The virtual appliance scratch disk.

6   Optionally, enter a new password for the virtual appliance in the **VA Password** text box, and then re-enter it in the **Confirm password** text box.

7   If you want to perform file-level recovery from backups of Linux virtual machines, select **Use virtual appliance for Linux File Level Restore** to configure a VA to use for Linux file-level recovery.

If you select this option for a different VA in the future, that VA is used for Linux file-level recovery.

Process of configuring virtual appliance networking

If you have deployed one or more virtual appliances using a static IP configuration, and you need to change that configuration after deployment, you must do so from within the virtual appliance.

To configure Virtual Appliance networking, complete the following tasks:

1   Logging in to the virtual appliance

2   Changing the hostname

3   Changing the network configuration

4   Configuring DNS

Logging in to the virtual appliance

Complete the steps in the following procedure to log in to the virtual appliance.

To log in to the virtual appliance

1   From your VI Client, find the VA and power it on.

2   Using the console, login as **root** to the virtual appliance, using the default credentials below:

  • Username: root

  • Password: vzroot1

If you supplied your own password during the VA deployment, use that password.

Changing the hostname

Complete the steps in the following procedure to change hostname.

To change the hostname

1   Login to the virtual appliance.

2   Using the VI text editor, edit the network configuration file:

```
vi /etc/sysconfig/network
```

> **NOTE:** See Using the VI Editor for more information on using the VI editor.

3   Change the hostname to the new value. Save the file.

4   Restart the VM by typing **reboot** at the prompt.

Changing the network configuration

Optionally, you can complete the steps in the following procedure to change the network configuration.

To change the network configuration

1   Login to the virtual appliance.

2   Using the VI text editor, edit the configuration file for eth0:

    vi /etc/sysconfig/ifconfig.eth0

3   Edit the configuration file as needed. The text below is shown as an example:

    ONBOOT=yes
    SERVICE=ipv4-static
    IFACE=eth0
    IP=<static IP of VA>
    GATEWAY=<IP of gateway>
    PREFIX=<this is the subnet mask - your values may differ>
    BROADCAST=<your broadcast IP- the highest IP in the VA's subnet.>

4   Save the file.

5   Restart the VM by typing **reboot** at the prompt.

Configuring DNS

Depending on your environment, it may be necessary to configure the virtual appliance with DNS nameserver information. The steps to change an existing configuration are the same.

To configure DNS

1   Login to the virtual appliance.

2   Using the VI text editor, edit the *resolv.conf* file:

    vi /etc/resolv.conf

3   Enter the following text, substituting the appropriate values where indicated:

    search <domain1> <domain2> <domain3>
    nameserver <IP address>
    nameserver <IP address>

4   Save the file.

5   Restart the VA by typing **reboot** at the prompt.

Disabling DRS for the vRanger virtual appliance

If the vRanger appliance is configured on hosts that are part of a DRS enabled cluster, the VA(s) may be migrated to a different host due to DRS/HA activity. This causes replication failures as the virtual appliance is no longer in the expected location (as configured in vRanger). To prevent this disruption, it is important to disable DRS for the vRanger virtual appliances.

To disable DRS for the vRanger virtual appliance

1   In your VI client, find the cluster(s) containing the vRanger virtual appliances.

2   Right-click on the cluster, and select **Edit Settings**.

3   In the **Cluster Settings** dialog, select **Virtual Machine Options**.

Search for or find the vRanger virtual appliances. Right-click on each virtual appliance, and select **manual**.

Using vRanger

- Interface overview
- My Inventory view
- My Jobs view
- My Repositories view
- My Reports view

Interface overview

The vRanger interface consists of four views: My Inventory, My Repositories, My Jobs, and My Reports. Each view features the Main Toolbar as well as each pane's Toolbar.

In the left pane within these views, data displays nested in a tree structure. In general, the information that displays in the right pane depends on the node that you select in the left pane. Each node has a distinct icon, making navigation easy and hierarchy clear.

For instructions on using the Startup Wizard to configure the software for your environment, see Configuring vRanger through the Startup Wizard.
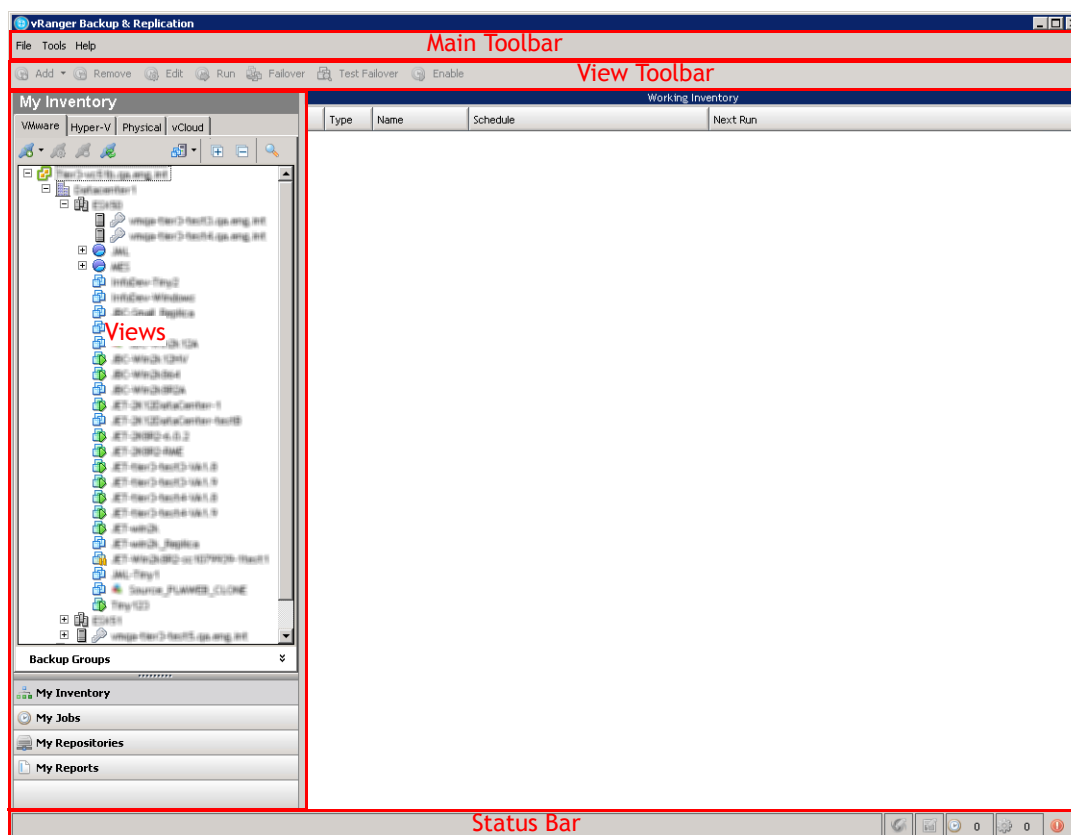
Main Toolbar

The Main Toolbar appears at the top of the vRanger user interface and includes the drop-down menus described in the following table.

File	The only submenu available under **File** is **Exit**. Click it to close the application.
Tools	• **Catalog Search:** Use this form to search and browse for savepoints that contain the files and folders you're looking for. Restore operations can be launched from the items in the results pane.
	• **Startup Wizard:** Select this command to launch the Startup Wizard, through which you can add a VC, host, repository; license a host; set up a mail server; and add email addresses.
	• **Options:** When you select this command, the Configuration Options dialog displays. It contains information in these categories: General, Virtual Appliance Configuration, Licensing, and My Jobs.
Help	• **Help Contents:** Select this command to view online help.
	• **License Information:** The License File dialog includes the total number of licensed CPUs, the total number of days that the license covers, and the number of days remaining under the current license. From this dialog, you can import a license file.
	• **About vRanger Backup & Replication:** Click this option to view version number, copyright information, a link to the website, and licensing terms. Also access the System Information dialog, which includes hardware resources, components, software environment, and Internet settings. The dialog features a window and submenus that let you view remote computers and system history. The About vRanger screen includes a System Info button that collects information about the vRanger machine. You may export this information to a text to aid interactions with Dell Support.

Status Bar

In every view, several icons display in the Status Bar at the bottom of the vRanger UI. Regardless of the view displayed, the Status Bar always gives you a glimpse of active job and task count. Using the example from the table below, the Active Jobs Count (4) corresponds to the first number in the My Jobs: Current Jobs node. The Active Tasks Count (2) corresponds to the second number at the same node level. That is, these counts display in this format in My Jobs: Current Jobs (4,2).

	Service Support Logging: Right-click to enable log files to assist in troubleshooting Windows vRanger service.
	Client Support Logging: Right-click to enable log files to assist in troubleshooting the vRanger client (GUI).
	Active Jobs Count: This indicator tells you the total number of active backup, restore, and replication jobs running in real-time.
	Active Tasks Count: This indicator tells you the total number of active backup, restore, and replication tasks running in real-time.
	New Failed Jobs: Indicates that there are new failed jobs.
	When you click the **New Failed Jobs** icon, this action resets the data in the Status Bar. In place of the new failed jobs icon, this icon displays and the active job and task counts are reset, populating with zeroes (0) initially.

My Inventory view

The My Inventory view displays the virtual infrastructure, configured jobs, and job details. It consists of the Inventory pane on the left and the Working Inventory pane on the right. This is the view from which most backup and replication jobs are added.
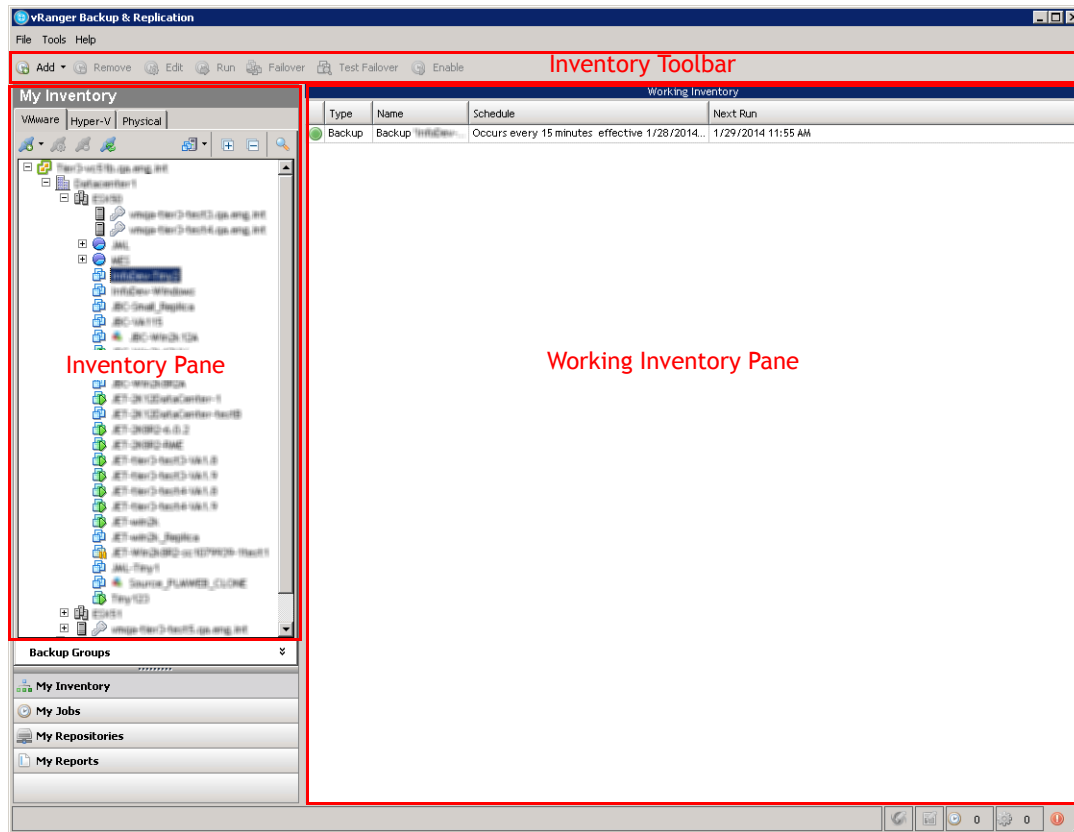
My Inventory pane

The My Inventory pane offers three tabs—one to show the infrastructure of VMware virtual machines, one to show the inventory of Hyper-V virtual machines, and one for physical servers.

From this pane, you can navigate and refresh the tree; add or remove connections to VCs or ESX Servers; and add credentials to virtual servers. You can also add a backup or replication job.

Inventory Pane toolbar and icons

The toolbar is located under the General Toolbar (File, Tools, Help). Options on this toolbar are enabled or disabled based on which object is selected.

Add: Click this command to add a backup or replication job.

Cancel: To cancel a job that is running, click this command.

Edit: Click this icon to open the appropriate wizard and make updates to the configuration for the job.

Disable: When you click this command, the job that is selected in the **Working Inventory** pane is disabled. A pause icon displays adjacent to the listing.

Enable: When a disabled job is selected in the **Working Inventory** pane, click this command to enable the job. A run icon displays next to the listing.

Remove: When you click to highlight a job or savepoint in the **Working Inventory** pane and then click this icon, the job is removed.

Run: Click this command to run the job that is selected. This reloads the window and displays the job in **My Jobs: Current Jobs**.

The icons below appear in the My Inventory Pane, and are enabled based on which tab you are viewing, and what you select in the **My Inventory** and **Working Inventory** panes
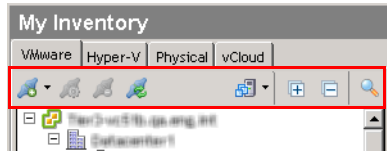
Add: The action this command triggers depends upon the node that you have highlighted in the **My Inventory** pane.

- **VMware tab** - click to add a Virtual Center, ESXi host, or backup group
- **Hyper-V tab** - click to add a Hyper-V System Center VMM, failover cluster, host, or backup group.
- **Physical tab** - click to add a physical machine or backup group.
- **vCloud tab** - click to add a vCloud Director server.

Properties: When a VC is selected, click Properties to view the credentials used to connect to it.

Remove: This command is available to delete any object that has been added to the view.

NOTE: If you remove an ESX host, all login and backup information is deleted. If there is a job scheduled to use this host, the job fails.

Refresh: Click this icon to refresh the My Inventory pane. Initially, only the highest-level node displays (VC). A Loading Inventory message displays while the pane refreshes.

Run: Click this icon to run an associated job.

Inventory: Click one of these icons to display the **Hosts and Clusters** or **Virtual Machines and Templates** view.

NOTE: This applies only to the VMware tab.

Expand All: This expands the tree view. After refreshing the pane, click this icon to reload the tree structure under the VC node.

Collapse All: Click this icon to collapse the tree view so that only the VC node displays in the pane.

Search: Click this icon to open the Search Inventory text box and find a specific item in the selected inventory pane.

VMware inventory tab

The VMware tab offers two views, depending on which value you select from the Inventory drop-down menu — Hosts and Clusters, or Virtual Machines and Templates.
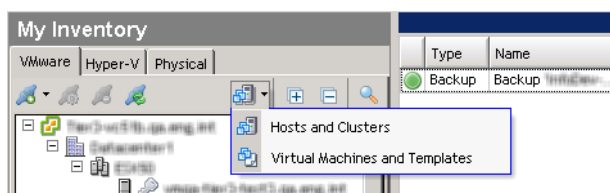
Hosts and clusters

When the Hosts and Clusters inventory option of the VMware tab is selected, the icons below appear in the inventory list, depending on the tab selected.

Folder: Click this icon to expand the view.

Resource Pool: Click this icon to expand the view

VirtualCenter: Right-click this icon to access these commands:
- Add Credential
- Credential Properties
- Remove VirtualCenter
- Backup VirtualCenter
- Restore
- View History
- Refresh Inventory

ESX(i) Cluster: Right-click this icon to access these commands:
- Add Credential
- Credential Properties
- Remove Credentials
- Backup
- Restore
- View History
- Refresh History

ESX(i) Host (not credentialed): Right-click this icon to access these commands:
- Add Credentials
- Credential Properties
- Remove Credentials
- Backup
- Restore
- Configure Change Tracking (choose any or all VMs in host)
- View History
- Refresh Inventory

ESX(i) Host (Credentialed): Right-click the icon to access these commands:
- Add Credentials
- Credential Properties
- Remove Credentials
- Backup
- Restore
- Configure Change Tracking (choose any or all VMs in host)
- View History
- Refresh Inventory

ESX(i) Host Add: Click this icon to add an ESX host.

ESX(i) Host Disconnected Right click this icon to connect to the host.

ESX(i) Host Error: Right click on this icon to determine the error.

ESX(i) Host Maintenance Mode: Indicates the host is in maintenance mode, and may not be available.

Hosts and Clusters

vApp: vApps can comprise of any applications running on any OS, and provide a mechanism for customers to move their applications between internal clouds or external clouds with still the same service levels.

Change Block Tracking: This icon indicates that a job has Change Block Tracking enabled. Click the icon to disable.

Virtual Machine Powered Off: Right-click this icon to access a menu of the following commands:
  - Add Host Credentials
  - Credential Properties
  - Remove Credentials
  - Backup [machine name]
  - Replicate [machine name]
  - Enable Change Tracking
  - View History
  - Refresh Inventory

Virtual Machine Powered On: Right-click this icon to access a menu of the following commands:
  - Add Host Credentials
  - Credential Properties
  - Remove Credentials
  - Backup [machine name]
  - Replicate [machine name]
  - Enable Change Tracking
  - Configure VSS Tools
  - View History
  - Refresh Inventory

Virtual Machine Suspended (Paused): Right-click this icon to access a menu of the following commands:
  - Add Host Credentials
  - Credential Properties
  - Remove Credentials
  - Backup [machine name]
  - Replicate [machine name]
  - Enable Change Tracking
  - View History
  - Refresh Inventory

Virtual machines and templates

When the Virtual Machines and Templates inventory option is selected, the icons below appear in the inventory list.

VirtualCenter: Right-click at this node level to access these commands:

- Add Credential
- Credential Properties
- Remove VirtualCenter
- Backup VirtualCenter
- Restore
- View History
- Refresh Inventory

Virtual Center Add: Click this icon to Add a Virtual Center.

Cluster: Right-click this node to access these commands:

- Add Credential
- Credential Properties
- Remove Credentials
- Backup Demo
- Restore
- View History
- Refresh History

Virtual Machine: Right-click the node to access these commands:

- Add Credential
- Credential Properties
- Remove Credentials
- Backup
- Restore
- Replication
- Enable (Disable) Change Tracking
- View History
- Refresh Inventory

Virtual Machine Folder: Click on this icon to display the virtual machines in this folder.

Virtual Machine Folder Open: Clicking this icon collapses the view and closes the folder.

Virtual Machine Powered Off: Right-click this icon to access a menu of the following commands:

- Add Host Credentials
- Credential Properties
- Remove Credentials
- Backup [machine name]
- Replicate [machine name]
- Enable Change Tracking
- View History
- Refresh Inventory

Virtual Machine Powered On: Right-click this icon to access a menu of the following commands:

- Add Host Credentials
- Credential Properties
- Remove Credentials
- Backup [machine name]
- Replicate [machine name]
- Enable Change Tracking
- Configure VSS Tools
- View History
- Refresh Inventory

Restore Virtual Machine: Clicking this icon starts a restore job for this VM.

Virtual Service: Click this icon to expand the list of services.

Backup Job: Right-click the node to access these commands:

- Add Credential
- Credential Properties
- Remove Credentials
- Backup
- Restore
- Replicate
- Enable Change Tracking
- View History
- Refresh Inventory

Running Backup Job: Right-click the node to access these commands:

- Add Credential
- Credential Properties
- Remove Credentials
- Backup
- Restore
- Replicate
- Enable Change Tracking
- View History
- Refresh Inventory

Backup Virtual Machine: Right-click on the preferred VM, then click this icon to start a backup job for this VM.

Replicate Virtual Machine: Right-click on the preferred VM, then click this icon to start a replication job for this VM.

Change Block Tracking: This icon indicates that a job has Change Block Tracking enabled. Left click the icon to disable.

Connected: This icon shows that the node is connected.

Connecting: This icon indicates that a node is in the process or connecting.

Credentialed: This icon indicates that VM or Host is Credentialed.

Data Center: Selecting this icon expands the data center and displays all objects in it.

Disconnected: This icon indicates the Node is disconnected.

Error: This icon indicates that an error has occurred, or that more data is required.

Hierarchy: This icon indicates the organization of the inventory displayed.

Inventory Tree Wait: This icon is displayed while vRanger searches for all the components in the inventory.

Credentialed/Licensed: This icon shows that the VM or host is licensed.

Invalid Licensed Credentials: This icon shows that the credentials are invalid.

Resource: This icon indicates the presence of a resource in the inventory.

Save: Clicking this icon saves any changes made.

Server: This icon indicates a server in the inventory.

Virtual Machine Template: copy and paste this template as needed to reproduce VMs.

vApp: vApps can comprise of any applications running on any OS, and provide a mechanism for customers to move their applications between internal clouds or external clouds with still the same service levels.

BackupReplication License: Indicates the VC is licensed for both backup AND replication.

Backup or **Replication License:** Indicates the VC is licensed for either backup OR replication.

Hyper-V inventory tab

When the Hyper-V tab is selected, the following icons are used to indicate the type and status of Hyper-V items.

Hyper-V Cluster: Right-click this icon to access the following commands:
- Credential Properties
- Remove From Inventory
- Reinstall Host Agent
- Uninstall Host Agent
- View History
- Backup [machine name]

NOTE: Uninstalling or reinstalling the host agent from a cluster node respectively removes or adds the agents to all nodes under that cluster.

Hyper-V Host: Right-click this icon to access the following commands:
- Credential Properties
- Remove From Inventory
- Reinstall Host Agent
- Uninstall Host Agent
- View History
- Backup [machine name]

Hyper-V Host Warning: This icon indicates that agent present on the host is not compatible with the version of vRanger in use.

Virtual Machine Powered Off: Right-click this icon to access the following commands:
- View History
- Backup [machine name]

VM Powered On: Right-click this icon to access the following commands:
- View History
- Backup [machine name]

Physical inventory tab

From the Physical tab, you can add or modify a physical server, or create a backup group consisting of multiple physical servers.

Physical Machine: Right-click this icon to access a menu of the following commands:
- Credential Properties
- Remove From Inventory
- Install (Reinstall) Host Agent
- Uninstall Host Agent
- Backup [machine name]
- View History
- Refresh Inventory

Physical Machine Warning: This icon indicates that agent present on the physical machine is not compatible with the version of vRanger in use.

Physical Machine Disconnected: This icon indicates that vRanger cannot communicate with the physical machine. Reasons for this may be but are not limited to: the machine is powered off, the agent is not installed on the machine, there is a network error, and so on.

Working Inventory pane

The Working Inventory pane lists any configured jobs for the node selected in the Inventory pane.

The data presented here - including Type, Name, and Schedule - can be sorted by dragging a column heading. When you click, an arrow displays in the column heading, indicating the sort order (ascending or descending). Click the column heading again to change the sort order. The same data can be grouped by dragging a column heading to the target area above all of the headings. Doing this groups current jobs by that data type.

If you click to highlight a job in the Working Inventory pane and then click the Run icon, the job is moved to the My Jobs: Scheduled Jobs. From this location, you can start the job. After the backup has been completed, it is moved to My Jobs: Recent Jobs.

My Jobs view

The My Jobs view consists of the My Jobs pane on the left, the Activity pane on the upper right, and the Job Detail pane on the lower right. Use this view to add, run, schedule, disable, or cancel a backup job. From this view you must have either the Scheduled Jobs or On Demand Jobs node (in the My Jobs pane) highlighted to add a job.
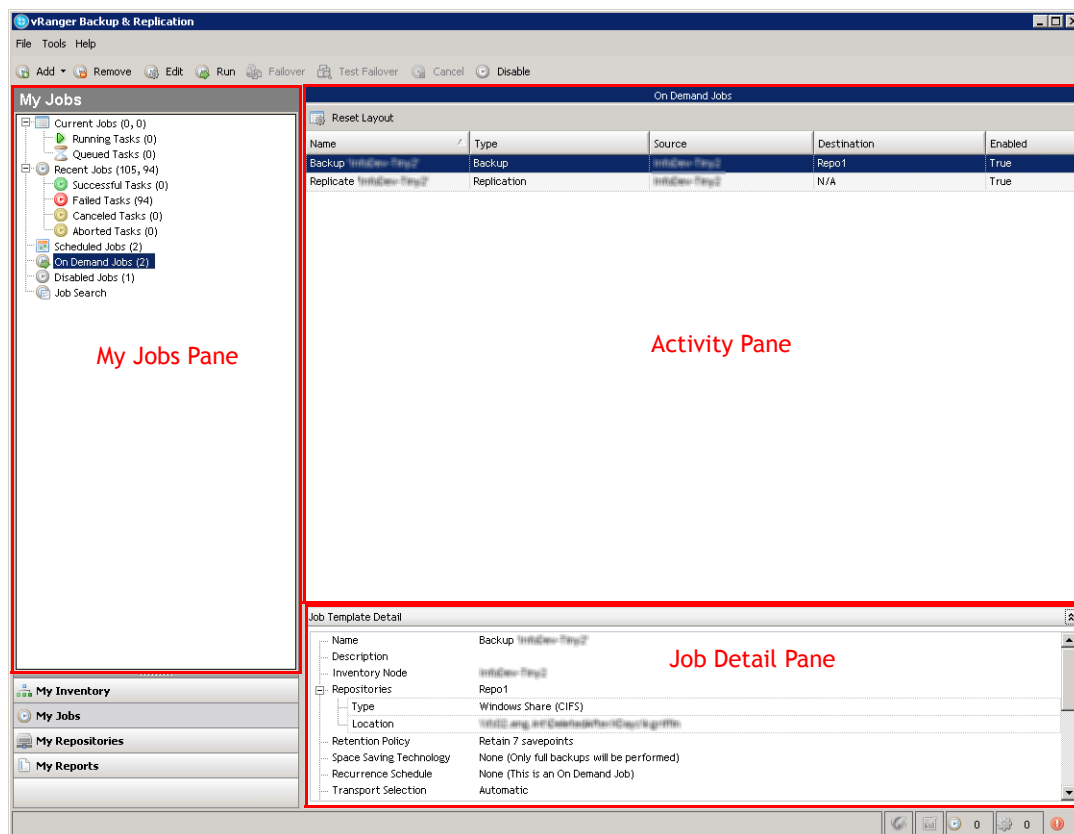
A job can be made up of many tasks. Even if some of its tasks fail, the job itself can complete successfully. In this view, you can also confirm configuration settings for a job. My Jobs makes monitoring jobs and tasks easy. Use the Job Detail pane to track job progress. To view task progress, see the Task Detail area. Use the tool to view job history and check status.

My Jobs pane

The My Jobs pane displays task and job information for each of the following statuses: Current; Recent; Scheduled; On Demand; and Disabled. For Current and Recent jobs, associated tasks are listed in separate categories. In these nodes, jobs and tasks are distinguished in this format —(x, y) —where *x* is the number of jobs and *y* refers to the number of tasks that make up the job. The numbers in the Current Jobs node corresponds to the values in the Status Bar's Active Jobs Count and Active Tasks Count areas, respectively.

My Jobs pane toolbar and icons

Add: Click this command to add a backup, restore, or replication job.

Remove: Use this icon to delete a job highlighted in the **Activity** pane.

Edit: With a job selected, click the this icon to access the **Backup Wizard**, **Restore Wizard**, or **Replication Wizard**.

Enable: When a disabled job is selected in the **Working Inventory** pane, click this command to enable the job. A run icon displays next to the listing.

Disable: When you click this command, the job that is selected in the Working Inventory pane is disabled. A pause icon displays adjacent to the listing.

Run: Click this command to run the job that is selected. This reloads the window and displays the job in **My Jobs: Current Jobs**.

Cancel: To cancel a job that is running, click this command.

Current Jobs: Click this icon to display Current Jobs. Two levels can be displayed:
- Running Tasks
- Queued Tasks

Running Tasks: Refers to the number of active tasks associated with current jobs.

Queued Tasks: For all current jobs, this number indicates the number of tasks waiting to run.

Recent Jobs: This node lists the number of recent jobs and tasks, respectively, that have been run. The format (*x*, *y*) is used, indicating *x* recent jobs and *y* recent tasks. Under this node, there are several objects:
- Successful Tasks
- Failed Tasks
- Canceled Tasks
- Aborted Tasks

Successful Tasks: This category indicates the number of tasks that have completed successfully for recent jobs.

Failed Tasks: The number in this category notes the failed tasks associated with recent jobs.

Canceled Tasks or **Aborted Tasks:** These tasks use the same icon but refer to different statuses for recent jobs. A canceled task results from a user action. An aborted task refers to a change initiated by the system.

Scheduled Jobs: When this node is selected, the **Activity** pane is divided into two tabs — **List View** and **Calendar View**. The jobs and tasks listed on both views are those that have been scheduled to run. The **Job Template Detail** area features configuration options selected for each job.

On Demand Jobs: These jobs are configured to run at any time; that is, they are not scheduled. The **Job Template Detail** area features configuration settings for each job.

Disabled Jobs: Listed here are jobs that have been configured, have been scheduled and perhaps even run, but are currently disabled. From this view, any disabled job can be enabled.

Selecting this node displays **Job** and **Results** areas in the **Activity** pane. A **Job Template Detail** area displays as well. Enter criteria in the Job area and view the data returned in the Results area. If you click a listing in the results, you can view additional detail about it in the **Job Template Detail** area.

LAN Failover: If the Failover option (along with the Fiber Backups option) was selected, this icon indicates the Fiber connect failed, and the job failed over to a LAN connection.

Activity pane

- The data that displays in the Activity pane depends upon the category that you select in the My Jobs pane. You can expand a job listing to display data on tasks associated with that job. For the Recent Jobs, On Demand Jobs, and Disabled Jobs, the Activity pane features data organized in these categories:
- **Name** – This is the name of the job.
- **Type** – This refers to job type.
- **Source** – Indicates the host or cluster in which the source server resides.
- **Destination** – For backup jobs, this indicates the repository selected in the backup job. For restore and replication jobs, this indicates the target host or server.
- **Status** – Lists the status of the backup or restore job or task: Completed, Failed, Aborted, or None.
- **Start Time** – Lists when the job or task was initiated.
- **Duration** – The length of time that the job or task took to complete.

- **Message** – In this column, information about a task or job displays. For example, "Host is not licensed" or "Please check host credentials."

If you select the Current Jobs node, all of the categories listed above display in the Activity pane. In addition, data for the following categories is included: Throughput and Progress.

In the interface, you can clearly distinguish the running job (parent) and the running task (child). To view details for either job or task, click to highlight the object. Either the Job Detail or Task Detail area displays at the bottom of the Activity pane.

You can sort the columns in the Activity pane by clicking them. When you click, an arrow displays in the column heading, indicating the sort order (ascending or descending). Click the column heading again to change the sort order.

If you select any sub-category under the Recent Jobs node and then right-click on a job or task in the Activity pane, a Go To menu displays. The available commands list depends upon the activities previously associated with that job or task. For example:

- Go To Repository Savepoint

- Go To Scheduled Activity

- Go To Disabled Activity

If you select a category under Current Jobs, Scheduled Jobs, or On Demand Jobs and then right-click on a job or task in the Activity pane, the following navigation options display:

- Go To Inventory View

- Find Previous Jobs

- Go To Current Jobs

For Scheduled Jobs, the Activity pane features two tabbed areas — List View and Calendar View. List View displays scheduled jobs by data type — Name, Type, Source, Destination, Enabled, and Next Run. The Calendar View offers monthly, weekly, daily options for viewing scheduled jobs. Below this, the Job Template Detail area features basic configuration information about the job selected in the List View.

Click the node in the My Jobs pane to display Job and Results areas in the Activity pane. After the results display, you can click to highlight any entry and additional information about the job displays in the Job Template Detail area below the results, which you can save.

Type	• Historic • On Demand • Scheduled • Disabled
Task Status	• All • Success • Failed • Aborted • Canceled
Repository	The values included here depend on the repositories that were added previously.
Job Type	• All • Backup • Restore • Replication
VM Name	To search by VM name, populate this text box.

Job Status	• All
	• Success
	• Failed
	• Aborted
	• Canceled
To Date	The default value is (none). When you click the drop-down, a calendar view displays.
Host or VC	The default value is (All). This is populated with the hosts and VCs that were added previously.

Job Detail area

The **Job Detail** area contains job configuration data as well as the progress and status of currently running jobs and tasks. This area features two tabs — Detail and Log. For example, if you select Recent Jobs in the My Jobs pane, the Job Detail area displays this type of data on the Detail tab. The data in the corresponding Log tab indicates which VMs or servers were retrieved, which disks were skipped, and which tasks were completed. It also indicates that the job was saved.

My Repositories view

This view displays the configured repositories and available savepoints. This is the view from which most restore jobs are configured. It consists of the My Repositories pane on the left and the Working Repository pane on the right.

My Repositories pane

This pane displays the configured repositories in a tree structure, grouped by repository type. Repositories can be added and removed through the pane toolbar.

My Repositories pane toolbar and icons

Catalog: After selecting a repository, click this icon to search (or browse) for a specific Savepoint.
NOTE: Cataloging must have been enabled at the time the backup job ran.

File-Level Restore: When you want to restore an individual file from a savepoint, click this icon to access the **File Level Recovery** dialog.

FLR From Manifest: When you want to restore an individual file from the Manifest, click this icon.

Restore: When you click to highlight a backup job or savepoint in the **Working Repository** pane, click this icon to open the **Restore Wizard**.

Remove: To delete a savepoint that is selected in the **Working Repository** pane, click this command.

Add: Use this command to add either a CIFS or SFTP repository.

Properties: This icon is the gateway to the connection information and repository credentials used to connect to the share. You can update credentials at any time.

Remove: Click this icon to open the **Repository Removal Confirmation** dialog. Before removing a repository, confirm that it is not needed to store current jobs.

Working Repository pane

The Working Repository pane shows the configured repositories and the savepoints stored there. The data displayed in this pane can be grouped and sorted using the icons in the Group By drop-down list. The data in this pane is organized under these columns:

- Repository Name
- VM Name
- Host Name
- Type (Full, Differential, Incremental)
- Start Time
- Original Size (MB)
- Stored Size (MB)

Grouping

There are two methods for grouping savepoints. You can select an option from the Group By toolbar, or you can create customized groupings using the GroupBox command.

Repositories & VMs: Groups the savepoints first by repository, then by VM.

Repositories: Click this icon to group savepoints by repository.

VMs: This command groups the savepoints by VM.

Date Range [Last 3 Months ▾] **Savepoints:** This groups the savepoints first by repository, then VM, then by root savepoint ID.

Hosts: This groups the savepoints by host.

Date Range [Last 3 Months ▾] **Date Range:** This drop-down list lets you select a date range for the jobs listed in the Working Repository pane. The values include:

- All
- Last Day
- Last 3 Days
- Last Week
- Last 2 Weeks
- Last Month
- Last 3 Months
- Custom
- Define Custom.

If you select Define Custom, the Savepoints Date Ranger dialog displays. Use it to set a date range for the jobs you want to view.

Sorting

The columns in the Working Repository pane can be re-ordered by dragging a heading and clicking the column. When you click, an arrow displays in the column heading, indicating the sort order (ascending or descending). Click the column heading again to change the sort order.

My Reports view

From the My Reports view you can create and manage reports as well as email them. The left-hand Reports Pane displays a list of all reports for backups and restores. The right-hand pane, called Working Reports, displays the report itself.
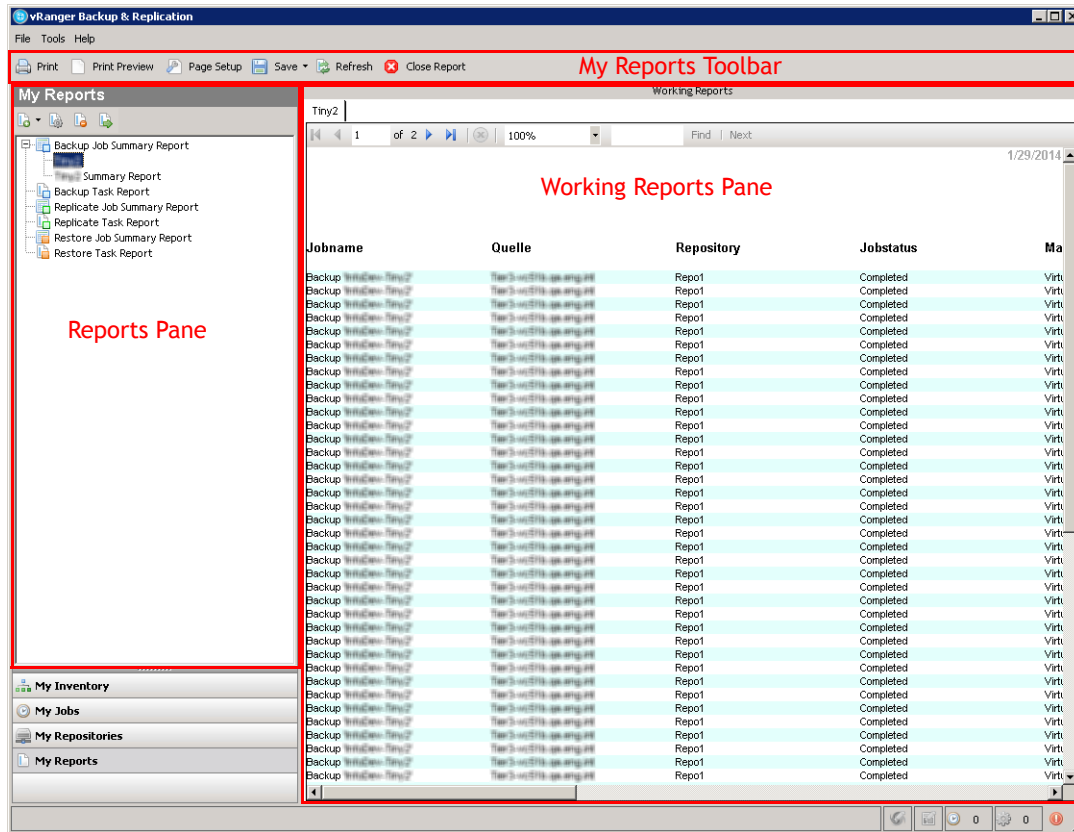
My Reports toolbar

The toolbar is located under the General Toolbar (File, Tools, Help). The options on this menu are only enabled if a report is showing in the Working Reports Pane, and act only on that report

Reports pane

Print: Use this icon to print the selected report.

Print Preview: Use this icon to preview the report before printing.

Page Setup: Use this icon to configure the page settings used to print your report.

Save: Use this icon to save your report as a PDF or Excel file.

Refresh: Use this icon to refresh the selected report.

Close: Use this icon to close the selected report.

The Reports Pane includes nodes for the following types of reports:

- Backup Job Summary

- Backup Task Summary

- Backup Task

- Replicate Job Summary

- Replicate Task

- Restore Job Summary

- Restore Task

You can also create a new report or work with existing reports through the Reports Pane toolbar icons. For more information, see Reports chapter.

Reports pane toolbar

Add: Use this icon to start the process of creating a backup or restore report.

Edit: Click this icon to open the **Report Wizard** to make changes to the selected report's configuration. Enabled only when a job is selected in My Jobs.

Remove: Click this icon to delete the selected report. Enabled only when a job is selected in **My Jobs**.

Run: Click this icon to run the selected report. Enabled only when a job is selected in **My Jobs**.

Backup

Backup overview

vRanger offers two types of backups - virtual and physical. Each type is described below. Where information relates to only one type of backup, that restriction is noted in the section title or as a note in the text.

Virtual machine backup

vRanger protects two varieties of virtual machines:

- VMware
- Hyper-V

Refer to the sections below for more information on the different types of virtual backups.

Protecting VMware virtual machines

When protecting a VMware virtual environments, vRanger uses VMware snapshot technology to temporarily store incoming write requests while the source VMs are being backed up. After a VMware backup completes, the snapshot is deleted, which commits those pending writes to disk. vRanger can backup a VMware VM that already has an open snapshot and can backup the open snapshot, but any secondary consolidated helper snapshots are closed prior to running the backup.

With Inventory Node Selection, you can browse the vCenter or vCloud Director inventory and select which VMs, groups, or vApps you want to protect. You can select a VM, folder, resource pool, vApp, ESX host, datacenter, or vCenter and backup all of the VMs located under that node in the tree.

Protecting VMware vCloud Director vApps

VMware vCloud Director vApps are containers - similar to a resource pool - that provide the ability to manage and control resources for the VMs contained within. vApps are generally used to properly group the individual components of a multi-tier application, allowing for application-level resource control and portability. When protecting vApps, vRanger backs up both the virtual machines and the vApp metadata to maintain the vApp settings and structure upon restore.

Protecting Hyper-V virtual machines

When protecting a Hyper-V host, the vRanger Hyper-V Agent is installed on the host and works with the VSS writer present on all Hyper-V hosts to back up every VM located on the host and send the backups directly to the repository. If the host is located on the same network as the repository location, such as a SAN, then the client can perform LAN-free backups.

ⓘ | **IMPORTANT:** Backing up a Hyper-V virtual machine using a VSS writer requires a minimum of 300 MB free space on the volume being backed up. If there is less than 300 MB of available space on the volume, the backup job fails with an error message.

VM backup limitations

vRanger cannot backup physical RDM (raw device map) partitions. You should see the warning: Incompatible drive detected.

Physical machine backup

For backing up physical machines, vRanger uses a physical backup client, similar to the one used for protecting Hyper-V VMs, to perform backup and communication operations. The client can be installed when the source server is added to the vRanger inventory, or installed manually to comply with change control requirements.

After it is installed, the physical client does the following:

- Communicates with the vRanger server
- Gathers information on the source server configuration
- Manages and executes the vRanger backup tools and operations
- Manages vRanger VSS snapshot activity on the source server
- Transmits backup data directly to the repository

vRanger physical backups are performed using the direct-to-target architecture. Depending on how the source server is connected to the repository, this can result in network backups or LAN-Free backups. Note that in order for backups to be LAN-Free, the source server and repository must exist on the same SAN.

ⓘ | **IMPORTANT:** Physical backups are only supported on Windows machines matching the list of supported operating systems. Refer to "*System requirements and compatibility*" in the *Dell vRanger Pro Installation/Upgrade Guide*.

VMware backup transport options

vRanger offers different data transport options for backup tasks. When configuring a backup job, vRanger automatically chooses the fastest option available, but you can also override this selection by configuring a custom transport selection. vRanger chooses from the following transport options, based on your configuration.

ⓘ | **NOTE:** The backup transport method describes only how data is read from the source server, not how the data is sent to the repository.

HotAdd

vRanger includes support for VMware's HotAdd disk transport functionality, which permits LAN-Free backups from a virtual machine.

vRanger can use HotAdd for two backup methods:

- **VA-based HotAdd** – Mounts the source VM's disk to the vRanger Virtual Appliance deployed on the source host (or cluster). This lets vRanger have direct access to the VM data through VMware's I/O stack rather than the network.

- **Machine based HotAdd** – If vRanger is installed in a virtual machine, this method mounts the source VM's disk to the vRanger virtual machine. This lets vRanger have direct access to the VM data through VMware's I/O stack rather than the network. With this method, the backup processing activity occurs on the vRanger server.

To use vRanger with HotAdd, the virtual machine (either the vRanger VA, or the vRanger virtual machine) must be able to access the target VM's datastore or datastores. In addition, all hosts to which the vRanger VM or VA could be vMotioned must be able to see the storage for all VMs that vRanger is configured to back up.

> ⓘ **NOTE:** When using HotAdd, plan on one concurrent backup per CPU core. To calculate the maximum number of concurrent backup tasks per proxy server, simply identify the number of CPU cores on that server - that is the maximum number of concurrent backups. For example, a Dual-Socket, Quad-Core system can perform up to eight concurrent backup jobs.

Configuring the vRanger machine for HotAdd

When using HotAdd, make sure to disable Automount on the vRanger machine. This prevents Windows on the vRanger VM from assigning a drive letter to the target VMDK. To disable Automount:

- From the start menu, select "run" and enter **diskpart**.

- Run the **automount disable** command to disable automatic drive letter assignment.

- Run the **automount scrub** command to clean any registry entries pertaining to previously mounted volumes.

HotAdd limitations

The use of HotAdd for backups is subject to the following limitations:

- **VMFS Block Size** – To use HotAdd, the VMFS containing the source VM and the VMFS containing the vRanger VM must use the same block size. For example, if you back up a virtual disk on a datastore with 1MB blocks, the proxy must also be on a datastore with 1MB blocks. For more information, refer to VMware's *VDDK 1.2.1 Release Notes*: http://www.vmware.com/support/developer/vddk/VDDK-1.2.1-Relnotes.html#knownissues

- **HotAdd and the vRanger VM** – You cannot use HotAdd to back up the vRanger virtual machine. This may result in the vRanger volume being removed from the VM along with the source VMDK. For more information, refer to VMware's *VDDK 1.2.1 Release Notes*: http://www.vmware.com/support/developer/vddk/VDDK-1.2.1-Relnotes.html#knownissues

- IDE disks are not supported.

- Mixed mode clusters are not supported for HotAdd. For example, a VM on an ESXi 4.0 host cannot HotAdd a disk from an ESX 3.5 host.

For information on HotAdd best practices and limitations, refer to VMware KB article 1035096.

SAN backups

To perform SAN backups, vRanger must be installed on a physical system attached to your fibre or iSCSI network. In addition, the volumes containing the servers to be protected must also be properly zoned/mapped to the vRanger proxy server.

Configuring vRanger for SAN backups

With vRanger installed on a physical proxy server, the following configurations must be made:

- Disable automount on the vRanger machine:

  From the start menu, click **Run**, and then enter **diskpart**.

  Run the **automount disable** command to disable automatic drive letter assignment.

- Verify that the SAN policy is set to Online All by typing **san** and hitting Enter.

  If it is not, set it to online all by typing **san policy=onlineAll**.

- Run the **automount scrub** command to clean any registry entries pertaining to previously mounted volumes.

- On your storage device, zone your LUNs so that the vRanger HBA (or iSCSI initiator) can see and read them.

- Only one proxy should see a set of LUNs at one time. For backups only, The proxy server should only have **read-only** access to the LUNs. To perform LAN-Free restores, ensure that the vRanger server has Read + Write access to any zoned LUNs to which you want to restore.

LAN

If no LAN-Free option is available, backups can be processed over the network. LAN backups are supported in all configurations and require no additional set-up to work, but are generally the slowest option. Network backups can be processed one of two ways, depending on the source host configuration.

- **VA-based LAN** – Transfers the source VM's data from the source disk to the vRanger Virtual Appliance (VA) over the network. With this method, the backup processing activity occurs on the VA.

- **Machine-based LAN** – If there is no vRanger VA deployed, vRanger transfers the source VM's data from the source disk to the vRanger server over the network. With this method, the backup processing activity occurs on the vRanger server. For ESXi servers (which do not have the service console), data is sent by way of VMware's VDDK transport. For ESX servers that are credentialed at the host level, Machine-based LAN backups are processed in the host Service Console.

ⓘ | **NOTE:** Generally speaking, network backups and restores are slower than LAN-Free. If your environment supports LAN-Free backups, use that option whenever possible.

Transport selection order - virtual machine backup

When backing up virtual machines or other virtual objects (hosts, clusters, and so on), the vRanger Backup Wizard includes an **Automatic** option which lets vRanger select the best available method for your configuration. When determining the best transport, vRanger uses two key criteria:

- Whether virtual appliances are deployed
- Where vRanger is installed

The logic used to select a transport differs slightly based on whether vRanger is installed in a VM or on a physical server. The sections below describe the steps taken to select an automatic transport.

vRanger - installed on a virtual machine

When installed on a virtual machine, vRanger checks the possible transports in the order shown below, selecting the first option available to ensure the best performance possible.

1 **Virtual Appliance on source host** – vRanger first checks the source host for a virtual appliance. If available, that VA is used.

   Is VA-based HotAdd available? – vRanger checks the local VA for HotAdd support. If the host is not properly licensed, or the VA cannot access the storage for the source VM, HotAdd is not available.

   ⓘ | **IMPORTANT:** If a virtual appliance is configured and HotAdd is not available, a network backup is performed from the virtual appliance. No further transport checks are performed.

2 **Virtual Appliance on source cluster** – If the source host does not have a virtual appliance, vRanger checks if the host is a member of a cluster, If so, vRanger checks for a shared virtual appliance on that cluster.

   Is VA-based HotAdd available? – vRanger checks the cluster VA for HotAdd support. If the host is not properly licensed, or the VA cannot access the storage for the source VM, HotAdd is not available.

   > ⓘ **IMPORTANT:** If a virtual appliance is configured and HotAdd is not available, a network backup is performed from the virtual appliance. No further transport checks are performed.

3 **Machine-based-HotAdd** – If there is no virtual appliance configured, vRanger performs a check to see if HotAdd is available for the vRanger VM. If the vRanger VM is on a host which is not properly licensed, or if the VM cannot access the storage for the source VM, HotAdd is not available.

4 **Machine-based-LAN** – If a virtual appliance is not available, and HotAdd is not supported on the vRanger VM, a network backup is performed from the vRanger machine.

vRanger - installed on a physical machine

When installed on a physical machine, vRanger checks the possible transports in the order shown below, selecting the first option available. The preferred transports are checked first, ensuring the best performance possible.

1 **Virtual Appliance on source host** – vRanger first checks the source host for a virtual appliance. If available, that VA is used.

   Is HotAdd available? – vRanger checks the local VA for HotAdd support. If the host is not properly licensed, or the VA cannot access the storage for the source VM, HotAdd is not available.

   > ⓘ **IMPORTANT:** If a virtual appliance is configured and HotAdd is not available, a network backup is performed from the virtual appliance. No further transport checks are performed.

2 **Virtual Appliance on source cluster** – If the source host does not have a virtual appliance, vRanger checks if the host is a member of a cluster, If so, vRanger checks for a shared virtual appliance on that cluster.

   Is HotAdd available? – vRanger checks the cluster VA for HotAdd support. If the host is not properly licensed, or the VA cannot access the storage for the source VM, HotAdd is not available.

   > ⓘ **IMPORTANT:** If a virtual appliance is configured and HotAdd is not available, a network backup is performed from the virtual appliance. No further transport checks are performed.

3 **Machine-based SAN** – If there is no virtual appliance configured, vRanger performs a check to see the vRanger server is configured for SAN backups.

4 **Machine-based LAN** – If a virtual appliance is not available, and SAN backups are not supported on the vRanger server, a network backup is performed from the vRanger machine.

Hyper-V backup transport options

When backing up Hyper-V virtual machines, backup activity is processed by the vRanger Hyper-V Agent on the Hyper-V host. Backup data is read by the agent and sent from the source server directly to the repository.

When performing backups of Hyper-V VMs, or backup groups containing Hyper-V VMs, the Transport Selection portion of the Backup Wizard is not shown, as there is only one available option.

Physical backup transport options

vRanger transports describe only how data is read from the source server, not how the data is sent to the repository. When backing up physical servers, the vRanger client reads the data directly from the disk, and sends data from the source server to the repository. Depending on how the source server is connected to the repository, this can result in network backups or LAN-Free backups. Note that in order for backups to be LAN-Free, the source server and repository must exist on the same SAN.

When performing backups of physical servers, or backup groups containing physical servers, the Transport Selection portion of the Backup Wizard is not shown, as there is only one available option.

Determining application consistency

The levels of, and methods for providing, application consistency differ between virtual machine backups and backups of physical servers. See the sections below for further details:

- Application consistency for virtual backups
- Application consistency for physical backups

Application consistency for virtual backups

vRanger provides various levels of protection for virtual machines. The means each level uses to accomplish its task depends on the type of virtual machine being protected.

VMware virtual machines

vRanger, by default, provides no quiescing during backups. When enabled (via the **Enable Guest Quiescing** option), quiescing in vRanger is provided by leveraging VMware Tools installed in the VM.

This can provide three different levels of backup consistency, as described below:

- **Crash Consistent** – A crash-consistent backup is analogous to pulling the plug on a server and then backing up the data. The state of the data that is being backed up with respect to the users of the data is indeterminate. Restoring a crash-consistent image is equivalent to rebooting a server after a hard shut-down.

- **File System Consistent** – File-system consistency is achieved through standard quiescing (via the VMware Sync Driver), which ensures that no file system writes are pending when the snapshot is taken. For normal VMs, file-system consistency is adequate, although it can cause corruption in database applications.

- **Application Consistent** – Consistency of VSS-compatible applications is achieved by freezing application I/Os just prior to creating the VM snapshots. This ensures that all application writes requests in the machines memory are committed to disk before the snapshot is taken.

The level of consistency provided by the **Enable Guest Quiescing** option is dependent upon the version of VMware ESX(i) (and the corresponding VMware Tools) and the guest operating system. The following table provides more detail on what is needed to achieve various levels of consistency:

Table 3. Achieving consistency

File-level quiescing

ESX version	Windows Server 2003	Windows Server 2008 (includes R2)	Windows Server 2012
ESX 5.0	VMware VSS	VMware VSS	VMware VSS
ESX 5.1	VMware VSS	VMware VSS	VMware VSS

Table 3. Achieving consistency

ESX 5.5	VMware VSS	VMware VSS	VMware VSS
ESX 6.0	VMware VSS	VMware VSS	VMware VSS
Application-level quiescing			
ESX version	**Windows Server 2003**	**Windows Server 2008 (includes R2)**	**Windows Server 2012**
ESX 5.0	VMware VSS	vzShadow.exe	vzShadow.exe
ESX 5.1	VMware VSS	vzShadow.exe	vzShadow.exe
ESX 5.5	VMware VSS	vzShadow.exe	vzShadow.exe
ESX 6.0	VMware VSS	vzShadow.exe	vzShadow.exe

As shown in the table above, application consistency is not always available with the basic quiescing options. In these situations, you may use vRanger's VSS Tools (vzShadow.exe) for application-level consistency.

Hyper-V virtual machines

vRanger can provide two levels of consistency for Hyper-V VM backups: file-system consistency and application consistency.

- **Crash Consistent** – A crash-consistent backup is analogous to pulling the plug on a server and then backing up the data. The state of the data that is being backed up with respect to the users of the data is indeterminate. Restoring a crash-consistent image is equivalent to rebooting a server after a hard shut-down.

- **Application Consistent** – vRanger engages Microsoft's Volume Shadow Copy Services (VSS) to put supported applications into a consistent state during a backup. This ensures that all application writes requests in the machines memory are committed to disk before the snapshot is taken, which means that the application and data can be reliably recovered from the backup archive.

To achieve either type of application consistency, you must comply with the following requirements:

- The guest must be operating a Windows operating system.
- The guest operating system must have Hyper-V Integration Services installed.

The vRanger agent installed on the Hyper-V host uses the VSS native to the Windows operating system to support the truncation of any supported application transaction logs. Log truncation automatically frees space in the logical log for the transaction log to reuse.

> ⓘ **IMPORTANT:** Backing up a Hyper-V virtual machine using a VSS writer requires a minimum of 300 MB free space on the volume being backed up. If there is less than 300 MB of available space on the volume, the backup job fails with an error message.

> ⓘ **NOTE:** For more information on Hyper-V backups and VSS, please refer to Microsoft's article Backing Up and Restoring Virtual Machines.

Application consistency for physical backups

vRanger can provide two levels of consistency for physical backups: crash consistent and application consistency.

- **File System Consistent** – File system consistency is achieved without any additional tools or configurations. File-system consistency ensures that no file system writes were lost during the backup process.

- **Application Consistent** – vRanger engages Microsoft's Volume Shadow Copy Services (VSS) to put supported applications into a consistent state during a backup. This ensures that all application writes

requests in the machines memory are committed to disk before the snapshot is taken, which means that the application and data can be reliably recovered from the backup archive.

For physical backups, vRanger also supports (via Microsoft's VSS) the truncation of any supported application transaction logs when the VSS snapshot is creating. Log truncation automatically frees space in the logical log for reuse by the transaction log.

Understanding retention policies and space-saving technologies

vRanger retention policies define the minimum number of restorable savepoints to keep in a repository for a backup job. The actual number of savepoints retained at any one time depends on the type of backups being performed (full, differential, or incremental) and where in the full-differential or full-incremental cycle the retention check is being performed.

> ⓘ | NOTE: vRanger does not retire an expired savepoint if that savepoint is required to restore a differential or incremental savepoint still within the retention period.

A key concept in discussing retention policies is the notion of a "backup set." A backup set is a number of savepoints that are grouped together for two reasons: The first is to implement the retention policy; and the second is to maintain the parent-child relationships that exist when we choose to use differential or incremental backups. The backup set for each type of backup is different, as described below.

- **Full backups** – Each savepoint is a complete backup set. Deleting a savepoint has no bearing on the recoverability of any other full savepoint.

- **Incremental backups** – Incremental backup jobs backup only the blocks that have changed since the last backup - full or incremental. Restoring an incremental savepoint requires the parent full *and* every incremental between the full and the selected incremental.

- **Differential backups** – A differential backup contains the data that has changed since the last full backup. Each differential backup includes the contents of the previous differential, which means, in terms of retention, that only the parent full and the most recent differential are required for a restore.

Full backups

A backup job consisting only of full backups is the simplest case for retention policy configuration. Each savepoint is independent of the other savepoints, and can be retired without affecting any other savepoints. In other words, a backup set for full backups is equal to one savepoint - only one savepoint is required for a restore.

The scenario below depicts a daily full backup job, with a Savepoint Count of 7:

- **Day 1** - A full backup is taken.

- **Day 2** - A full backup is taken.

- **Day 3** - A full backup is taken.

- **Day 4** - A full backup is taken.

- **Day 5** - A full backup is taken.

- **Day 6** - A full backup is taken.

- **Day 7** - A full backup is taken.

- **Day 8** - A full backup is taken. The savepoint from Day 1 is removed.

- **Day 9** - A full backup is taken. The savepoint from Day 2 is removed.

Incremental backups

Incremental backups job back up only the blocks that have changed since the last backup - whether it was full or incremental. Restoring an incremental savepoint requires the parent full *and* every incremental between the full and the selected incremental

The scenario below depicts a daily incremental backup job, with a Savepoint Count of 7 and a Threshold Count of 6:

Table 4. Incremental backup retention example

Day	Action
1	A full backup is taken.
2	An incremental backup is taken.
3	An incremental backup is taken.
4	An incremental backup is taken.
5	An incremental backup is taken.
6	An incremental backup is taken.
7	An incremental backup is taken.
8	A full backup is taken.
9	An incremental backup is taken.
10	An incremental backup is taken.
11	An incremental backup is taken.
12	An incremental backup is taken.
13	An incremental backup is taken.
14	An incremental backup is taken. All savepoints from Day 1 to Day 7 are removed.
15	A full backup is taken.

Differential backups

Differential savepoints require the parent full backup and the selected differential savepoint to restore. Each differential backup includes the contents of the previous differential, which means, in terms of retention, that only the parent full and the most recent differential are required for a restore.

The scenario below depicts a daily differential backup job, with a Savepoint Count of 7, a Threshold Count of 6, and a Threshold Size of 50%:

Table 5. Differential backup retention example

Day	Action
1	A full backup is taken.
2	A differential backup is taken.
3	A differential backup is taken.
4	A differential backup is taken.
5	A differential backup is taken.
6	A differential backup is taken.
7	A differential backup is taken.
8	A full backup is taken. The differential savepoint from Day 2 is removed.
9	A differential backup is taken. The differential savepoint from Day 3 is removed.
10	A differential backup is taken.The differential savepoint from Day 4 is removed.
11	A differential backup is taken.The differential savepoint from Day 5 is removed.
12	A differential backup is taken.The differential savepoint from Day 6 is removed.

Table 5. Differential backup retention example

Day	Action
13	A differential backup is taken.
14	A differential backup is taken. The differential savepoint from Day 7 and the full savepoint from Day 1 are removed.

Performing optional configurations

Before you begin setting up a backup job, you may want to make the following optional vRanger configurations:

- Enabling or disabling VMware Changed Block Tracking
- Adding a custom backup group
- Configuring VSS tools for VMware virtual machines

Enabling or disabling VMware Changed Block Tracking

VMware Changed Block Tracking (CBT) reduces the time needed for incremental and differential backups by only backing up the portions of a disk that have changed since the last backup. By determining which blocks changed within the VMDK file, vRanger only backs up the portions of a disk that have changed since the last backup. This often results in shorter backup operations, and reduced resource consumption on network and storage elements.

vSphere supports CBT, and most VMs running in this environment can use it. The VMs must be Hardware Version 7 or above, and have been created and hosted in ESX4 or 4i hosts. VMs that are created in ESX3.5 or below must be migrated to Hardware Version 7 or above for CBT to be supported. CBT must be enabled for each VM with which CBT is to be used.

ⓘ | **NOTE:** Virtual machines that have not yet been powered on cannot have CBT enabled. You must have powered on the VM for the CBT option to be available.

To enable or disable VMware Changed Block Tracking

- On the **My Inventory** view, right-click the VM for which want to change the CBT setting, and then select one of the following options:
  - **Enable Tracking**
  - **Disable Tracking**

The **Tracking Icon** appears next to any VM on which CBT is enabled.

Adding a custom backup group

This feature lets a user organize physical servers, Hyper-V servers, or vCenter objects (VMs, hosts, and folders) into groups for the purpose of creating a backup job for only the selected nodes. This provides additional flexibility for creating backup jobs.

ⓘ | **IMPORTANT:** A backup group must contain only physical objects, only Hyper-V objects, or only VMware objects. You cannot combine VMware, Hyper-V, and physical objects into a single backup group.

To add a custom backup group

1 On the My Inventory view, select the VMware, Hyper-V, or Physical tab.

2 Do one of the following:

- Click the **Add** Icon and select **Backup Group.**

- Or, do the following:

    - Expand the Backup Group panel in the Inventory pane.

    - Right-click an empty spot, select **Add**, and then click **Backup Group**.

3   Enter a name for the **Backup Group**, and click to select the nodes to include in the backup group.

4   Clicking **Add** closes the **Add Backup Group** dialog — your new Backup Group is listed in the **Backup Groups** pane.

You can edit or remove a group by right-clicking it and selecting **Edit** or **Remove**.

Configuring VSS tools for VMware virtual machines

VSS Tools are only available for VMware virtual machines (VMs) and do not apply to Hyper-V VMs or physical machines.

There are two ways to configure VSS Tools: automatically and manually. The following sections describe the procedures for the two available methods of configuring VSS Tools with vRanger.

Configuring VSS tools using the automated method

In addition to the standard quiescing implementation using VMware Tools, vRanger provides an optional method for application quiescing on VMware virtual machines, which uses a small executable tool (*vzShadow.exe*) deployed to the source VM and activated by VMware Tools during quiescing. vRanger's VSS Tools must be deployed to each VM for which you need application-consistent backups. Refer to Application consistency for virtual backups for more information.

> ⓘ | **NOTE:** This method is available only for powered-on VMs running WIndows 2008 (or newer). For VMs running older Windows operating systems, you need to install vzShadow.exe manually, as described in Configuring VSS Tools Using the Manual Method.

To configure VSS tools using the automated method

1   When viewing the **My Inventory View**, right-click the preferred VM and select **Configure VSS Tools**.

2   The **Configure VSS Tools** dialog appears.Enter the name and password for an account with administrative privileges on the VM. Click **Next**.

3   If the selected VM has more than one disk, the available disks are shown in the **Disk for VSS snapshot section**. Select the disks to quiesce.

4   If preferred, select **Perform application log truncation**.

5   Click **Configure**.

6   Click **Next**.

vRanger performs the following actions:

- Create the directory *C:\Program Files\Dell\vRanger\VSS*

    populate the directory with the *vzShadow* executable and supporting files.

- Create the directory *C:\Program Files\VMware\VMware Tools\backupscripts.d*

    create a freeze.bat file in that directory that contains the appropriate contents based on your selections.

Removing VSS tools

To remove the vRanger VSS Tools and batch file from the source VM, complete the steps in the following procedure.

To remove VSS tools

1. When viewing the **My Inventory View**, right-click the preferred VM and select **Configure VSS Tools**.

2. The **Configure VSS Tools** dialog appears.Enter the name and password for an account with administrative privileges on the VM. Click **Next**.

3. Click **Remove**. The files are removed from the VM, but the directories remain.

Configuring VSS tools using the manual method

For Windows 2003 VMs, the vRanger VSS Tools must be manually deployed to the target VM, using the procedures below.

ⓘ | **NOTE:** vzShadow.exe works in conjunction with VMware Tools. VMware Tools must be installed and running on the VM when using vzShadow.exe.

To configure VSS tools using the manual method

1. On the vRanger machine, browse to the "*Client\VSS\VSSfor2k3*" sub-directory in the vRanger installation location. If the default location is accepted during the installation, the path is: *C:\Program Files\Dell\vRanger\Client\VSS\VSSfor2k3*

2. In this directory, find the appropriate C++ Redistributable package and *vzShadow.exe* for the bit level of the source VM (the one you want to back up).

   For 64-bit (x64) machines, choose

   vcredist_x64.exe (Microsoft C++ Redistributable)

   vzshadow_x64.exe

   For 32-bit (x86) machines, choose

   vcredist_x86.exe

   vzshadow_x86.exe

3. On the source VM, launch the *vesicatories* file to install the C++ Redistributable.

4. Depending on your system, move the appropriate file to the machine's root directory (typically the C:\ directory).

   For 64-bit OS, copy *vzshadow_x64.exe*

   For 32-bit OS, copy *vzshadow_x86.exe*

5. Create a directory on the VM:

   C:\Program Files\VMware\VMware Tools\backupscripts.d

6. In this directory, create a batch file called **freeze.bat**.

   A sample batch file is shown below;

   ```
   c:\vzshadow.exe x:
   ```

   where X equals the drive that hosts the application — if there are multiple volumes on this VM, list all volumes, separated by colon space (:), as in:

   ```
   c:\vzshadow.exe c: d: e:
   ```

When using vRanger with Dell Recovery Manager for Exchange 5.1 (or later), use the text below for your batch file. This creates an XML file that RME uses to identify vRanger savepoints with Exchange databases.

```
c:\vzshadow_64.exe -bc=x:\backup.xml x:
```

```
c:\vzshadow_86.exe -bc=x:\backup.xml x:
```

where X equals the drive that hosts the Exchange Server — if there are multiple volumes on this VM, list all volumes, separated by colon and a space (:), as in:

```
c:\vzshadow_x86.exe -bc=c:\backup.xml c: d: e:
```

When VSS is triggered by VMware tools, the pre-freeze script executes *vzShadow.exe*.

Configuring a backup job

vRanger uses a wizard to guide you through the process of creating a backup job. Each page of the wizard is explained in more detail in the sections that follow.

To configure a backup job, complete the following tasks:

1 Selecting an inventory node

2 Excluding virtual machines

3 Including hard disks of physical or virtual machines

4 Selecting a repository

5 Selecting a backup data transport method for VMware machines

6 Configuring backup options

7 Selecting a Retention Policy

8 Setting the recurrence schedule

9 Configuring email notifications

10 Starting the backup job

Selecting an inventory node

The first step to creating a vRanger backup job is to select the inventory node (physical server, vCenter, host, Backup Group, and so on) that you want to back up. After selecting the node, you may launch the Backup Wizard as described in the procedures below.

You may also create a custom grouping consisting of multiple physical servers or of existing objects from multiple hosts or vCenters. See Adding a custom backup group for more information. Backup Groups are treated the same as any other inventory node when it comes to creating backup jobs.

> ⓘ **IMPORTANT:** A backup group must contain only physical objects, only Hyper-V objects, or only VMware objects. You cannot combine VMware, Hyper-V, and physical objects into a single backup group.

To select an inventory node, complete one of the following procedures:

- Selecting a node in the My Inventory view

- Selecting a node in the My Jobs view

> ⓘ **NOTE:** If you want to back up more than one node, you need to first create a Custom Backup Group. For more information, see Adding a Custom Backup Group.

Selecting a node in the My Inventory view

When working in the My Inventory view, launching the Backup Wizard creates a job to back up the highlighted inventory node or backup group. If a node s not highlighted, the Add icon in the toolbar is disabled.

To select a node in the My Inventory view

1 In the **My Inventory** view, click the appropriate tab (VMware, Hyper-V, Physical, or vCloud), and then select the preferred node to highlight the object or backup group.

   ⓘ | **NOTE:** When viewing the vCloud inventory, the only inventory object that can be selected for backup is a vApp.

2 Do one of the following:

   • Right-click the highlighted node and select **Backup [object name]**.

   • Or, with the preferred inventory node selected, click the **Add** drop-down menu, and then click **Backup Job**.

   The **Backup Wizard** launches.

3 In the Job Name text box, enter a name for the backup job.

4 Optionally, in the Job Description text box, enter a description for the backup job.

5 Click **Next**.

Selecting a node in the My Jobs view

You may also launch the Backup Wizard from the Scheduled Jobs, On Demand Jobs, and Disabled Jobs node of the My Jobs view. As there should be no inventory object selected, the process is slightly different than when using the My Inventory View.

To select a node in the My Jobs view

1 In the **My Jobs** view, click one of the following nodes to enable the Add drop-down menu:

   • Scheduled Jobs

   • On Demand Jobs

   • Disabled Jobs

   ⓘ | **NOTE:** The node you select has no bearing on the status of the created job. For example, selecting the Disabled Jobs node does not cause the job to be created in a disabled state.

2 Click the **Add** drop-down menu, and then click **Backup Job**.

   The Backup Wizard opens to the Inventory Node Selection page.

3 Click the VMware, Hyper-V, Physical or vCloud tab.

4 Select the node you want to back up.

   • To back up a custom backup group, expand the Backup Groups pane, and then select the group of nodes displayed.

5 In the Job Name text box, enter a name for the backup job.

6 Optionally, in the Job Description text box, enter a description for the backup job.

7 Click **Next.**

Excluding virtual machines

If you selected a node that includes multiple servers (such as a backup group or host) and multiple virtual machines (VMs), the next page in the Backup Wizard is the **Virtual Machine Exclusion** page, followed by the Hard Disk Inclusion page. If you selected a single VM or are backing up physical servers, you see only the Hard Disk Inclusion page.

This page of the wizard only displays if you created a job for a node that includes multiple virtual machines, such as a host or backup group. On this page, you can select individual VMs to exclude from the backup job. For

example, if you want to back up all but two of the VMs on a host, you can configure the backup job for the host node, and then use this page to select the two VMs you do not want to include in the job.

To exclude virtual machines

1   In the **Exclude** column on the left, select the VMs you want to exclude from the backup job.

2   Click **Next**.

On the Virtual Machine Hard Disk Inclusion page, only the VMs included in the backup job appear in the Current Virtual Machines pane.

Including hard disks of physical or virtual machines

The Hard Disk Inclusion page follows the Virtual Machine Exclusion page. If you selected a single VM or are backing up physical servers, you see only the Hard Disk Inclusion page.

The Machine Hard Disk Inclusion page creates a rule for vRanger to use when determining which disks to back up.

When you create a backup job for one server, the rule tells vRanger which disks to back up for that server. When you create a job for multiple servers, however, the rule determines which disks are backed up for *every* server.

When protecting a virtual machine, the Current Virtual Machines pane displays the servers included in the backup job, based on your selections in the **Excluding virtual machines** procedure. The Include Hard Disks pane represents the maximum number of disks configured for any server in the job. For example, if you have a backup job with multiple single-disk servers and one server with six disks, the Include Hard Disks column shows six entries. Selecting "Hard Disk 1" through "Hard Disk 6" tells vRanger to back up every disk for each server.

If, in the example above, the six-disk server receives an additional disk, that disk is not going to be backed up, because only six disks are included in the rule. To allow for future growth, the **Show all disk possibilities** option lets you to select all potential disks, to ensure that any disk added to a server in the future is included in the backup job.

To include hard disks of physical or virtual machines

1   In the **Include Hard Disks** pane, select the disks that vRanger should include in the rule it applies to each server.

2   When including hard disks for VMs, the selection is not server-specific. It is global for all VM servers in this backup job. When including hard disks for physical machines, you can deselect disks for the backup group. However, you can not deselect Disk 1 for the group. Optionally, if you want to allow for additional disks in the future, select **Show all disk possibilities**, and select additional disks.

3   Click **Next**.

Selecting a repository

On the Repository Selection page of the wizard, you can select the repository to which the job should send the backup data. You can also add a new repository from this page. For more information, see Adding a repository.

vRanger supports the following repository options for storing backup archives:

- Windows Share (CIFS)

- SFTP

- FTP

- NFS

- **NetVault SmartDisk** - Dell Software's disk-based data-deduplication option which reduces storage costs with byte-level, variable-block-based software deduplication. For more information on NetVault SmartDisk, see http://software.dell.com/products/netvault-smartdisk/ or the *Dell vRanger Integration Guide for NetVault SmartDisk*.

- **Dell Rapid Data Access (RDA)** - Provided by the Dell DR Series appliances - purpose-built, disk backup appliances that use Dell deduplication technology to significantly improve backup and recovery processes. For more information on Dell DR Series appliances, see http://software.dell.com/products/dr-series-disk-backup-appliances/ or the *Dell vRanger Integration Guide for Dell DR Series Disk Backup Appliance*.

- **EMC Data Domain Boost (DDB)** - Integrating EMC® Data Domain® Boost™ with vRanger is achieved by adding a Data Domain appliance running DD Boost to vRanger as a repository. Backups written to that repository are deduplicated according to your configuration. For more information on EMC Data Domain Boost, see http://www.emc.com/data-protection/data-domain/data-domain-boost.htm or the *Dell vRanger Integration Guide for EMC® Data Domain® Boost™*.

ⓘ **NOTE:** Dell RDA and EMC DDB repositories support repository replication. For more information on configuring repository replication, refer to Managing Repository Replication.

ⓘ **NOTE:** Only one repository can be configured per backup job.

To select a repository

1   In the Repositories pane, select the repository you want to use for this backup job.

- If you want to add a new repository, see Adding a repository for more information.

ⓘ **NOTE:** When using Dell Rapid Data Access (RDA) or EMC Data Domain Boost repositories, you may configure the repository to replicate savepoints to another like device. For more information, refer to Managing Repository Replication.

2   Click **Next**.

Selecting a backup data transport method for VMware machines

In vRanger, the transport determines how backup data is sent, and where the backup processing activity occurs (for more information, see VMware backup transport options). When configuring a backup job for a VMware VM, the Backup Wizard, by default, uses the Automatic Transport Selection option to automatically select the best transport method at run-time, but you can also configure the transport manually by using the Using the Custom Transport Selection option.

ⓘ **NOTE:** The Transport Selection page is not available when backing up a physical machine or Hyper-V virtual server. Because only one transport method is available for Hyper-V and physical servers, the Transport Selection page of the Backup Wizard does not display.

To select backup data transport method for VMware Machines, complete one of the following procedures:

- Using the Automatic Transport Selection
- Using the Custom Transport Selection

Using the Automatic Transport Selection

The Backup Wizard includes an **Automatic Transport Selection** option, which lets vRanger select the best available method for your configuration. When determining the best transport, vRanger uses two key criteria:

- Whether virtual appliances (VAs) are deployed.
- Where vRanger is installed.

The logic used to select a transport differs slightly based on whether vRanger is installed on a VM or on a physical server. For more information about the order in which vRanger checks for transport availability, see Transport selection order - virtual machine backup.

To use the Automatic Transport selection

1   On the Transport Selection page, select **Automatic transport selection**.

    The selected transport method displays at the bottom of the page.

    ⓘ | **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method may not always be what is shown on the page during the job setup.

    vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

    • If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

        For more information about configuring the VA, see Understanding the vRanger virtual appliance.

2   Click **Next**.

Using the Custom Transport Selection

The Custom Transport Selection option lets you specify both where the actual backup processing activity occurs) and which transport to use. For more information, see VMware backup transport options.

To use the Custom Transport selection

1   On the Transport Selection page, select **Custom transport selection**.

    This enables the custom settings.

2   Select where you want the job to occur from the following options:

    • On this vRanger machine

    • On a vRanger Virtual Appliance

    vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

    • If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

        For more information about configuring the VA, see Understanding the vRanger virtual appliance.

        ⓘ | **NOTE:** Using the vRanger Virtual Appliance is recommended for scalability and performance.

3   Select which transports the job should attempt from the following options:

    • Advanced [SAN / HotAdd]

    • LAN

    For more information, see VMware backup transport options.

4   [Optional] If you want your LAN backups to be encrypted, select **Use SSL encryption while transferring files over the LAN connection when available**.

⚠ | **WARNING: This option causes reduced performance.**

5   Click **Next**.

Configuring backup options

The Backup Options page lets you configure additional operations performed during the backup job. The options available depend on the server, or inventory, type you want to back up.

To configure backup options

1   On the Backup Options page of the wizard, select your preferences from the options described in the following table.

Table 6. Available options

Option	Description	Inventory type
Backup powered on machines only.	If this flag is selected, vRanger checks the status of the VM before it begins a backup. If the VM is not running, it may not have changed significantly since the last backup, making another backup redundant.	• VMware • Hyper-V
Check destination for free space.	This option prompts vRanger to check the destination location to ensure that there is enough space available for the backup file to be written. Because it is not possible to determine how much space is saved by compression, vRanger generates an error if the free space at the destination is less than the combined size of the files belonging to the source server.	• VMware • Hyper-V • Physical
Compress backed up files.	This option prompts vRanger to compress the archive. **NOTE:** The following repositories do not support compression: NetVault SmartDisk, Data Domain DD Boost, and Dell Rapid Data Access (RDA).	• VMware • Hyper-V • Physical
Update notes with latest backup results.	If this flag is selected, vRanger updates the VM Notes in vCenter with the most recent backup status. vRanger maintains the existing VM notes in vCenter, appending the backup status to the end.	• VMware
Enable guest quiescing.	This option enables the VMware Tools when creating a backup. Use this flag if you are backing up a database server. During this process, vRanger freezes or pauses writes on supported applications to provide a consistent image of VSS-enabled databases. See Application consistency for virtual backups for more information.	• VMware
Enable VSS Application-level quiescing	This option engages Microsoft's VSS to put supported applications into a consistent state during a backup. **NOTE:** This option is only presented for physical backups. For backups of VMware and Hyper-V virtual machines, selecting **Enable guest quiescing** will enable VSS application quiescing if supported by the guest OS and application.	• Physical
Enable VSS log truncation	This option uses Microsoft's VSS to initiate the truncation of any supported application transaction logs upon a successful backup. **NOTE:** This option is only presented for physical backups. For backups of VMware and Hyper-V virtual machines, selecting **Enable guest quiescing** will enable VSS log truncation if supported by the guest OS and application.	• Physical

Option	Description	Inventory type
Enable Active Block Mapping (ABM)	Checking this box enables Active Block Mapping, which lets vRanger skip zeroed out blocks and junk data. **NOTE:** With this option enabled, deleted data on the source server is not backed up. The following ABM settings options are also available: • Do not include the Windows page file • Do not include the Windows hibernation file **NOTE:** In vRanger 6.1, page files were not backed up by default. In vRanger 7, this is no longer the case. If you are upgrading to vRanger 7 from version 6.1, page files will be backed up unless you edit each job and explicitly select the option to skip page files.	• VMware • Hyper-V • Physical
Enable Cataloging	Cataloging lets you browse, search, and restore file level content in a backup image. For more information, see Enabling cataloging options.	• VMware • Hyper-V • Physical

2   Click **Next**.

Selecting a Retention Policy

The Retention Policy Selection page lets you configure the type of backup to perform (full, incremental, or differential) and the minimum number of savepoints to keep for each VM. For more information about retention policies, see Understanding retention policies and space-saving technologies.

ⓘ | **IMPORTANT:** For backup jobs configured to use a replicated repository (see Managing Repository Replication), the chosen retention policy is applied to both the source and target repositories.

vRanger offers three methods of backing up data while reducing the storage footprint and network load of VM backups.

To create a backup job, complete one of the following procedures:

- Creating a full backup job
- Creating an incremental backup job
- Creating a differential backup job

Creating a full backup job

For full backups, vRanger performs a full scan and copy of the source VM. Full backup jobs can be modified with the following operations:

- **Active Block Mapping (ABM):** ABM scans the disks on a virtual machine and detects the blocks being actively used by the disk, as opposed to blocks that have been deleted by the Windows operating system. With ABM enabled, vRanger only backs up that part of a virtual disk that has active data on it. ABM also provides the following settings options:

  - Do not include the Windows page file

  - Do not include the Windows hibernation file

> **NOTE:** In vRanger 6.1, page files and hibernation files were not backed up by default. In vRanger 7, this is no longer the case. If you are upgrading to vRanger 7 from version 6.1, page files will be backed up unless you edit each job and explicitly select the option to skip the page file and/or hibernation file.
>
> ABM does not backup deleted data. If a VM backed up using ABM is restored, undelete operations are not possible.

To create a full backup job

1 On the Retention Policy Selection page, use the up and down arrows to select your Savepoint Count.

   The default Savepoint Count is 7.

2 In the Space Saving Technology section, select **None**.

3 Click **Next**.

Creating an incremental backup job

For incremental backups, vRanger performs a full scan of the source VM, but copy only blocks that have changed since the last backup.

Incremental backups are typically the fastest backup option and consume less storage space per archive. It is important to note, however, that to restore from an incremental backup, each incremental archive between the full backup and the preferred restore point must be available. When creating incremental backup schedules, use caution to minimize the risk introduced by long incremental chains.

Restoring an incremental backup can take longer than restoring a full or differential backup, because each previous incremental archive needs to be restored as well.

Incremental backups can be modified with the following operations:

- **Active Block Mapping (ABM):** ABM scans the disks on a virtual machine and detects the blocks being actively used by the disk, as opposed to blocks that have been deleted by the Windows operating system. With ABM enabled, vRanger only backs up the part of a virtual disk that has active data on it. ABM also provides the following settings options:

  - Do not include the Windows page file

  - Do not include the Windows hibernation file

    > **NOTE:** In vRanger 6.1, page files and hibernation files were not backed up by default. In vRanger 7, this is no longer the case. If you are upgrading to vRanger 7 from version 6.1, page files will be backed up unless you edit each job and explicitly select the option to skip the page file and/or hibernation file.

⚠ | **WARNING: ABM does not back up deleted data. If a VM backed up using ABM is restored, undelete operations are not possible.**

- **Changed Block Tracking (CBT):** Changed Block Tracking reduces the time needed for incremental and differential backups by only backing up the portions of a disk that have changed since the last backup. By determining which blocks changed within the VMDK file, vRanger only backs up the portions of a disk that have changed since the last backup. This often results in shorter durations for backup operations, and reduced resource consumption on network and storage elements.

  > **NOTE:** CBT does copy deleted blocks if ABM is not also enabled. CBT is a VMware feature, and does not apply to physical machine or Hyper-V virtual machine backup.

For more information about selecting the ABM and CBT operations, see Configuring backup options.

To create an incremental backup job

1 On the Retention Policy Selection page, select your Savepoint Count.

The default Savepoint Count is 7.

2   In the Space Saving Technology section, select **Incremental.**

The Options section appears.

3   Configure your incremental options by selecting one the following:

- **Threshold Count** – This defines the number of incremental backups to be performed before a full backup is done.

  Use the up and down arrows to select a threshold count. The default value is 6.

- **Full Backup Days** – A full backup is taken on the day or days you select.

  Select the day or days on which you want vRanger to take a full backup of the server.

  > ⓘ | **NOTE:** To run full backups only on the selected days, the Threshold Count must be set to a value of 8 or higher. A threshold count of 7 or less causes a full backup to also run seven days from the date the job is scheduled, even if that day is not a selected day.

4   Click **Next.**

Creating a differential backup job

For differential backups, vRanger performs a full scan of the source VM, but copies only blocks that have changed since the last Full backup.

Differential backups are typically slower than incrementals and can consume more storage space. They are usually faster to restore, however, and require only one differential backup and the parent full backup to restore.

Differential backups can be modified with the following operations:

- **Active Block Mapping (ABM):** ABM scans the disks on a virtual machine and detects the blocks being actively used by the disk, as opposed to blocks that have been deleted by the Windows operating system. With ABM enabled, vRanger only backs up the part of a virtual disk that has active data on it. ABM also provides the following settings options:

  - Do not include the Windows page file

  - Do not include the Windows hibernation file

  > ⓘ | **NOTE:** In vRanger 6.1, page files and hibernation files were not backed up by default. In vRanger 7, this is no longer the case. If you are upgrading to vRanger 7 from version 6.1, page files will be backed up unless you edit each job and explicitly select the option to skip the page file and/or hibernation file.

> ⚠ | **WARNING: ABM does not back up deleted data. If a VM backed up using ABM is restored, undelete operations are not possible.**

- **Changed Block Tracking (CBT):** Changed Block Tracking reduces the time needed for incremental and differential backups by only backing up the portions of a disk that have changed since the last backup. By determining which blocks changed within the VMDK file, vRanger only backs up the portions of a disk that have changed since the last backup. This often results in shorter durations for backup operations, and reduced resource consumption on network and storage elements.

  > ⓘ | **NOTE:** CBT copies deleted blocks if ABM is not also enabled. CBT is a VMware feature, and does not apply to physical machine or Hyper-V virtual machine backup.

For more information about selecting the ABM and CBT operations, see Configuring backup options.

To create a differential backup job

1   On the Retention Policy Selection page, select your Savepoint Count.

The default Savepoint Count is 7.

2   In the Space Saving Technology section, select **Differential.**

    The Options section appears.

3   Configure your incremental options by selecting one the following:

- **Threshold Count** – This defines the number of incremental backups to be performed before a full backup is done.

    Use the up and down arrows to select a threshold count. The default value is 6.

- **Full Backup Days** – A full backup is taken on the day or days you select.

    Select the day or days on which you want vRanger to take a full backup of the server.

    ⓘ | **NOTE:** To run full backups only on the selected days, the Threshold Count must be set to a value of 8 or higher. A threshold count of 7 or less causes a full backup to also run seven days from the date the job is scheduled, even if that day is not a selected day.

- The **Threshold Size** refers to the percentage of the hard disk of a VM that must be captured in differential backups before performing a full backup. Using the Threshold Size value may cause your full backups to be taken on unexpected days.

    ⓘ | **NOTE:** The Threshold Size setting is still applied if one or more Full Backup Days are set. To prevent the Threshold Size setting from triggering an unwanted full backup, set the Threshold Size to 100 or higher.

4   Click **Next**.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

ⓘ | **NOTE:** When using the Full Backup Days option, ensure that the selected recurrence schedule causes backup tasks to be run on the configured Full Backup Days.

To set the recurrence schedule

1   On the Recurrence Schedule Selection page, select from the options described in the following table.

This will be an on demand job ...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:

- **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
- **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
- **Monthly** – The monthly option offers the following configurations:
  - **Day *[x]* of every*[y]* month:**
    - *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
    - *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
  - **The *[f] [d]* of every *[y]* month(s):**
    - *f* can be either: first, second, third, fourth or last.
    - *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
- **Yearly** – The yearly option offers the following configurations:
  - **Every *[m][x]*:**
    - *m* is any month of the year. This determines the month of the year in which the report occurs.
    - *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
  - **The *[f][d]* of *[m]*:**
    - *f* can be either: first, second, third, fourth, or last.
    - *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      m is any month of the year. This determines the month of the year in which the report occurs.
- **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between backup jobs.

NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date:

- **No end date** – The job recurs until it is stopped.
- **End after** – Select the number of occurrences.
- **End by** – Select the date by which to end the recurrence.

2 Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

    • To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the backup job

The final page in the Backup Wizard displays a summary of the selections you made for the job.

To start the backup job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

    ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

    ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

    a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

    b   In the View toolbar, click **Run**.

    c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

Restore

- The vRanger restore process
- Process of performing a full restore for virtual machines
- Process of performing a full restore for VMware vCloud Director vApps
- Process of performing a full restore of a physical machine
- About file-level restore - Windows
- About performing a Linux file-level restore
- Restoring from manifest

The vRanger restore process

vRanger offers three main options for restoring from backup: restoring the full server from a repository; restoring a file or files from a repository; or restoring the server or file from a manifest. For more information, see the following sections:

- Full restore
- Restoring a file
- Restoring from Manifest

Full restore

Using vRanger, recovery of an entire VM (or resource pool, vApp, etc) is a simple process. A full restore returns the protected objects to the state as of the point in time at which the backup was performed. If the savepoint is for a differential backup, the restore process first restores the matched full backup, and then merge the data from the selected differential archive. For incremental savepoints, the full backup is restored along with each incremental archive taken between the full and the selected savepoint.

When restoring a parent object, such as a VMware vApp, vRanger restores the VM data as well as the parent object's metadata, allowing you to restore the entire object state and settings as well as the actual VM data.

During the normal restore process, information is pulled from the vRanger database to complete the restore job. In the event that the vRanger server is lost, and the database cannot be recovered (or is otherwise unavailable), you can reinstall a fresh version of vRanger and import the repository to regain access to your backups.

Restoring a physical server

vRanger restores physical savepoints to physical target servers using the vRanger Restore image (via CD or USB drive). This restore image provides a temporary boot environment into which the vRanger tools are loaded and executed. After the restore is complete, the server can be rebooted into the restored operating system.

After the target server is configured, the process of restoring a physical savepoint is similar to restoring a virtual machine. A full restore returns the server(s) to the state as of the point in time at which the backup was performed. If the savepoint is for a differential backup, the restore process first restores the matched full

backup, and then merge the data from the selected differential archive. For incremental savepoints, the full backup is restored along with each incremental archive taken between the full and the selected savepoint.

> ⓘ | **IMPORTANT:** When restoring to a physical server, ensure that you are restoring to hardware supported by the operating system and drivers in the savepoint being restored. Restoring a backup to dissimilar hardware may prevent the restored server from functioning.

See Process of performing a full restore of a physical machine for the procedures on restoring a physical savepoint.

Restoring a file

Unlike traditional backup solutions that require file-level agents, vRanger can recover a file or files directly from the image-level backup, without mounting the image. Incremental backups are combined with the parent full image as needed.

One of the key problems with file-level recovery is finding the appropriate files. vRanger includes a cataloging feature that indexes backups of Windows servers, and records the metadata to enable faster searching. For more information, see Performing a file-level restore using Catalog Search.

Restoring from Manifest

Each savepoint contains a manifest file that can be used to restore the savepoint without access to the vRanger database. In addition, the Restore from Manifest function can be used to recover savepoints not part of a repository (those that have been recovered from tape, for example).

The Restore from Manifest function can be used for full restores of VMs and physical servers, as well as file-level recovery. To restore a savepoint using the manifest file, see Restoring from manifest.

Process of performing a full restore for virtual machines

vRanger performs a full restore of a virtual machine (VM) by using a savepoint created from a backup job. A savepoint contains point-in-time data of the VM. If you select an incremental savepoint from which to restore the server, a full savepoint is also restored, along with every incremental savepoint taken between the time of the full backup and the time of the selected incremental backup.

For information about performing a full restore from a savepoint manifest, see Restoring from manifest.

For more information, see Process of performing a full restore of a physical machine.

To perform a full restore for a virtual machine, complete the following tasks:

1 Launching the Restore Wizard

2 Selecting a savepoint

3 Selecting a destination

4 Selecting a network

5 Selecting a restore data transport method for VMware machines

6 Configuring restore options

7 Setting the recurrence schedule

8 Configuring email notifications

9 Starting the restore job

Launching the Restore Wizard

When restoring a server from a savepoint, the Restore Wizard guides you through the steps necessary to complete the restore.

To use the Restore Wizard

1 In the main vRanger interface, click **My Repositories**.

2 Select the repository in which the savepoints for the server you want to restore are stored.

3 In the **Working Repository** pane, select a savepoint to restore.

4 Do one of the following:

- On the toolbar, click **Restore**.

- Or, right-click the savepoint, and then and click **Restore**.

   The vRanger Restore Wizard launches.

5 On the Mode Selection page of the Restore Wizard, enter the **Job Name** and **Job Description** in the appropriate text boxes.

6 Click **Next**.

Selecting a savepoint

The Savepoint Selection page displays all available savepoints, including their respective VMs and repositories. The savepoint selected on the Mode Selection page appears checked.

To select a savepoint

1 [Optional] Do one or more of the following:

- Clear the check box and select a different savepoint.

- Or, select an additional savepoint (for incremental backups).

   ⓘ | **NOTE:** If you select two or more savepoints with the same name on the same host, an error message displays and the Next button is disabled until you chose your selection.

2 Click **Next**.

Selecting a destination

The Destination Selection page lets you configure the host, drives, and datastores to which the savepoints are to be restored.

To select a destination

1 To select options to be applied to all VMs, do the following:

- Select the host from the **Master Host** drop-down list.

- Confirm that the path that appears in the Virtual Machine Location text box is correct.

   To change this path, use the browse (...) button to locate and a select a new location.

- Click the **Master Datastore** browse button to locate and select the datastore to be used for all VMs.

- Click **Assign**.

2 To select options for individual VMs, do the following:

- Select the disks you want to restore.

  Disks not selected appear in italics.

- To select a different host for each VM, click the drop-down list in the Restore To Host column, and then select the host to which the VM should be restored.

- To select a different datastore for each VM, browse and select a new location by doing one of the following:

  - For VMware VMs, click the button next to the **Primary Datastore** column entry.

  - For Hyper-V VMs, click the button next to the **Configuration File Location** column entry.

3 To rename to VM you are restoring, enter the new name under the Restore As column.

4 Click **Next**.

Selecting a network

The Network Selection page lets you configure the virtual network to which each vNIC should be connected upon restore. You can also configure a Master Network, which uses the selected network for all vNICs.

To select a network

1 To use the same network for all VMs, do the following:

   a Select a host from the **Master Network** drop-down list to which all VMs should be restored.

   b Click **Assign**.

2 To use different networks for individual VMs, select a virtual network from the connected network adapters included in the **Restore To Host** drop-down list for each VM to which the vNIC should be connected.

3 Click **Next**.

Selecting a restore data transport method for VMware machines

In vRanger, the transport determines how restore data is sent, and where the restore processing activity occurs (for more information, see VMware backup transport options). When configuring a restore job for a VMware VM, the Restore Wizard, by default, uses the Automatic Transport Selection option to automatically select the best transport method at run-time, but you can also configure the transport manually by using the Using the Custom Transport selection option.

ⓘ | **NOTE:** The Transport Selection page is not available when restoring a physical or Hyper-V virtual server. When restoring physical servers and Hyper-V virtual machines, vRanger sends the data to be restored over the network from the repository. Because this is the only method available for transporting data for physical and Hyper-V servers, the Transport Selection page of the Restore Wizard does not display.

To select a restore data transport method for VMware machines, complete one of the following procedures:

- Using the Automatic Transport selection

- Using the Custom Transport selection

Using the Automatic Transport selection

The Restore Wizard includes an Automatic Transport Selection option, which lets vRanger select the best available method for your configuration. When determining the best transport, vRanger uses two key criteria:

- Whether virtual appliances (VAs) are deployed.

- Where vRanger is installed.

The logic used to select a transport differs slightly based on whether vRanger is installed on a VM or on a physical server. For more information about the order in which vRanger checks for transport availability, see Transport selection order - virtual machine backup.

To use the Automatic Transport selection

1   On the Transport Selection page, select **Automatic transport selection**.

The selected transport method displays at the bottom of the page.

> ⓘ | **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method may not always be what is shown on the page during the job setup.

vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

- If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

    For more information about configuring the VA, see Understanding the vRanger virtual appliance.

2   Click **Next**.

Using the Custom Transport selection

The Custom Transport Selection option lets you specify both where the actual restore processing activity occurs) and which transport to use. For more information, see VMware backup transport options.

To use the Custom Transport selection

1   On the Transport Selection page, select **Custom transport selection**.

This enables the custom settings.

2   Select where you want the job to occur from the following options:

- On this vRanger machine.
- On a vRanger virtual appliance.

vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

- If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

    For more information about configuring the VA, see Understanding the vRanger virtual appliance.

    > ⓘ | **NOTE:** Using the vRanger virtual appliance is recommended for scalability and performance.

3   Select which transports the job should attempt from the following options:

- Advanced [SAN / HotAdd]
- LAN

For more information, see VMware backup transport options.

4   [Optional] If you want your LAN backups to be encrypted, select **Use SSL encryption when transferring files over the LAN connection**.

> ⚠ | **WARNING: This option causes reduced performance.**

5   Click **Next**.

Configuring restore options

The Options Selection page displays the restore options available depending on the type of VM you want to restore.

To configure restore options

1 On the Options Selection page, select from the following restore options:

- **Force overwrite** – This overwrites any existing disk or config files of the same name as those in the destination host.

- **Force power on** – This turns on the VM after the restore.

- **Create cluster role for restored virtual machine** – This option is for Hyper-V VMs only. It attempts to add the VM as a cluster resource of the Hyper-V host parent cluster.

  > **NOTE:** The cluster option is available only when a Hyper-V cluster has been added to the vRanger inventory and includes the host that is selected as the restore destination.

2 Click **Next**.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the restore job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

To set the recurrence schedule

1 On the Recurrence Schedule Selection page, select the options as described in the following table.

This will be an On Demand job…	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:

- **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
- **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
- **Monthly** – The monthly option offers the following configurations:
  - **Day *[x]* of every*[y]* month:**
    - *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
    - *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
  - **The *[f] [d]* of every *[y]* month(s):**
    - *f* can be either: first, second, third, fourth or last.
    - *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
- **Yearly** – The yearly option offers the following configurations:
  - **Every *[m][x]*:**
    - *m* is any month of the year. This determines the month of the year in which the report occurs.
    - *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
  - **The *[f][d]* of *[m]*:**
    - *f* can be either: first, second, third, fourth, or last.
    - *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      m is any month of the year. This determines the month of the year in which the report occurs.
- **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between jobs.

NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date:

- **No end date** – The job recurs until it is stopped.
- **End after** – Select the number of occurrences.
- **End by** – Select the date by which to end the recurrence.

2   Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

   • To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the restore job

The final page in the Restore Wizard displays a summary of the selections you made for the job.

To start the restore job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

   ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

   ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

   a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

   b   In the View toolbar, click **Run**.

   c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

Process of performing a full restore for VMware vCloud Director vApps

vRanger performs a full restore of a vApp by using a savepoint created from a backup job. A savepoint contains point-in-time data of the vApp, including member VMs and vApp metadata. If you select an incremental savepoint from which to restore the server, a full savepoint is also restored, along with every incremental savepoint taken between the time of the full backup and the time of the selected incremental backup.

For information about performing a full restore from a savepoint manifest, see Restoring from manifest.

For more information, see Process of performing a full restore of a physical machine.

To perform a full restore for a vApp, complete the following tasks:

1   Launching the Restore Wizard

2   Selecting a destination

3   Mapping Storage

4   Selecting a restore data transport method for VMware machines

5   Configuring restore options

6   Setting the recurrence schedule

Launching the Restore Wizard

When restoring a server from a savepoint, the Restore Wizard guides you through the steps necessary to complete the restore.

To use the Restore Wizard

1   In the main vRanger interface, click **My Repositories**.

2   Select the repository in which the savepoints for the server you want to restore are stored.

3   In the **Working Repository** pane, select a savepoint to restore.

4   Do one of the following:

   • On the toolbar, click **Restore**.

   • Or, right-click the savepoint, and then and click **Restore**.

   The vRanger Restore Wizard launches.

5   On the Mode Selection page of the Restore Wizard, enter the **Job Name** and **Job Description** in the appropriate text boxes.

6   Click **Next**.

Selecting a destination

The Destination Selection page lets you configure the virtual datacenter (vDC) to which the vApp savepoints are to be restored, as well as the name of the restored vApp.

To select a destination

1   To rename the vApp you are restoring, enter the new name in the **Restore vApp As** field.

2   In the **Restore vApp under vDC** drop-down, select the virtual datacenter to which the vApp should be restored.

3   Click **Next**.

Mapping Storage

The Restore Storage Mapping page lets you configure the storage profile and, as needed, the datastore for each VM in the vAPP.

To select storage mapping

1   In the **Storage Profile** drop-down, select the desired storage profile. Selecting a storage profile will allow vCloud Director to automatically select the appropriate datastore based on your storage profile settings.

2   To manually specify a datastore, select "Any" in the **Storage Profile** drop-down, then specify the desired datastore in the **Datastore** drop-down.

3   Click **Next**.

Selecting a restore data transport method

In vRanger, the transport determines how restore data is sent, and where the restore processing activity occurs (for more information, see VMware backup transport options). When configuring a restore job for a VMware VM or vApp, the Restore Wizard defaults to the Automatic Transport Selection option to automatically select the best transport method at run-time. You can also configure the transport manually by using the Using the Custom Transport selection option.

> ⓘ **NOTE:** The Transport Selection page is not available when restoring a physical or Hyper-V virtual server. When restoring physical servers and Hyper-V virtual machines, vRanger sends the data to be restored over the network from the repository. Because this is the only method available for transporting data for physical and Hyper-V servers, the Transport Selection page of the Restore Wizard does not display.

To select a restore data transport method for VMware machines, complete one of the following procedures:

- Using the Automatic Transport selection
- Using the Custom Transport selection

Using the Automatic Transport selection

The Restore Wizard includes an Automatic Transport Selection option, which lets vRanger select the best available method for your configuration. When determining the best transport, vRanger uses two key criteria:

- Whether virtual appliances (VAs) are deployed.
- Where vRanger is installed.

The logic used to select a transport differs slightly based on whether vRanger is installed on a VM or on a physical server. For more information about the order in which vRanger checks for transport availability, see Transport selection order - virtual machine backup.

To use the Automatic Transport selection

1   On the Transport Selection page, select **Automatic transport selection**.

    The selected transport method displays at the bottom of the page.

    > ⓘ **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method may not always be what is shown on the page during the job setup.

2   Click **Next**.

Using the Custom Transport selection

The Custom Transport Selection option lets you specify both where the actual restore processing activity occurs) and which transport to use. For more information, see VMware backup transport options.

To use the Custom Transport selection

1   On the Transport Selection page, select **Custom transport selection**.

    This enables the custom settings.

2   Select where you want the job to occur from the following options:

    - On this vRanger machine.
    - On a vRanger virtual appliance.

    vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

    - If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

For more information about configuring the VA, see Understanding the vRanger virtual appliance.

> ⓘ | **NOTE:** Using the vRanger virtual appliance is recommended for scalability and performance.

3   Select which transports the job should attempt from the following options:

- Advanced [SAN / HotAdd]
- LAN

For more information, see VMware backup transport options.

4   [Optional] If you want your LAN backups to be encrypted, select **Use SSL encryption when transferring files over the LAN connection**.

> ⚠ | **WARNING: This option causes reduced performance.**

5   Click **Next**.

Configuring restore options

The Options Selection page displays the restore options available for vApp restore jobs.

To configure restore options

1   On the Options Selection page, select from the following restore options:

- **Force overwrite** – This overwrites any existing disk or config files of the same name as those in the destination host.
- **Force power on** – This turns on the vApp after the restore.

2   Click **Next**.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the restore job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

To set the recurrence schedule

1   On the Recurrence Schedule Selection page, select the options as described in the following table.

This will be an On Demand job...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:

- **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
- **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
- **Monthly** – The monthly option offers the following configurations:
    - **Day *[x]* of every*[y]* month:**
        - *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
        - *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
    - **The *[f] [d]* of every *[y]* month(s):**
        - *f* can be either: first, second, third, fourth or last.
        - *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
        y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
- **Yearly** – The yearly option offers the following configurations:
    - **Every *[m][x]*:**
        - *m* is any month of the year. This determines the month of the year in which the report occurs.
        - *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
    - **The *[f][d]* of *[m]*:**
        - *f* can be either: first, second, third, fourth, or last.
        - *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
        m is any month of the year. This determines the month of the year in which the report occurs.
- **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between jobs.

NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date:

- **No end date** – The job recurs until it is stopped.
- **End after** – Select the number of occurrences.
- **End by** – Select the date by which to end the recurrence.

2 Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1 On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

    • To select all email addresses, select the check box in the Selected column header.

2 Click **Next**.

Starting the restore job

The final page in the Restore Wizard displays a summary of the selections you made for the job.

To start the restore job

1 On the Summary page, review the data displayed in the tree view.

2 If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

    ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3 Click **Finish**.

4 If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

    ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

    a In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

    b In the View toolbar, click **Run**.

    c In the Confirm Run Job dialog box, click **Yes**.

To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

Process of performing a full restore of a physical machine

Like the restore process for a virtual machine, the restore process for a physical machine uses the vRanger Job Wizard. While the process of restoring physical machines is similar to the virtual machine restore process, some of the preliminary preparations are different.

For information about performing a full restore from a savepoint manifest, see Restoring from manifest.

To perform a full restore of a physical machine, complete the following tasks:

1 Creating the boot media

2 Booting the target server

3 Selecting a savepoint

4 Naming the job

5 Mapping disks

6 Setting the recurrence schedule

7 Configuring email notifications

Creating the boot media

To restore a savepoint to a physical machine, you must first boot the target server into the vRanger restore environment using the vRanger Restore images (ISO for CD, IMG for a bootable USB drive).

When vRanger is installed, two restore image files (*vzRestore-1.2.6.iso* and *vzRestore-1.2.6.img*) are extracted to the "Service" sub-directory of the vRanger installation folder. By default this is on of the following:

- On x86 machines – *C:\Program Files (x86)\Dell\vRanger\Service*
- On x64 machines – *C:\Program Files\Dell\vRanger\Service*

To create the boot media, complete one of the following procedures:

- Creating a bootable CD
- Creating a bootable USB drive

Creating a bootable CD

Complete the steps in the following procedure to create a bootable CD.

> (i) | **IMPORTANT:** The procedures for creating a bootable CD differ widely depending on the operating system and software used. The procedures below show how to use the native Windows 7 utility to create a bootable CD. If you have a different image burning application, follow the instructions for that application.

> (i) | **NOTE:** This procedure requires a blank CD/DVD and a writable CD/DVD drive.

To create a bootable CD

1   Insert the blank CD/DVD into the writable drive.

2   In the vRanger installation directory, find the file *vzRestore-1.2.6.iso*. If necessary, copy this file to the machine with the writable CD/DVD drive.

3   Right-click on the ISO file, select **Open with**, and then click **Windows Disk Image Burner.**

4   From the disk burner drop-down list, select the drive containing the blank disk.

5   Click **Burn**.

Creating a bootable USB drive

Complete the steps in the following procedure to create a bootable USB drive.

> (i) | **NOTE:** There are no native Windows utilities for creating a bootable USB drive from an IMG file. The procedure below was created using a freeware application - Image Writer for Windows. Follow the instructions for your application.

To create a bootable USB drive

1   Insert a blank USB drive into your USB port.

2   In the vRanger installation directory, find the file *vzRestore-1.2.6.img*. If necessary, copy this file to the machine with the blank USB drive.

3   In the Image Writer for Windows UI, browse to the *vzRestore-1.2.6.img* file.

4   In the **Device** drop-down list, ensure that the correct USB drive is selected.

5   Click **Write**.

6   Click **Exit**.

Booting the target server

In order for vRanger to detect the target server, the target server must be booted into the vRanger restore environment and have a valid IP address. The vRanger Restore image is configured to use Dynamic Host Configuration Protocol (DHCP) by default. If you are using static IP, some additional steps are required.

To boot the target server, complete one of the following procedures:

- Booting into the restore environment using DHCP
- Booting into the restore environment using static IP

Booting into the restore environment using DHCP

Complete the steps in the following procedure to boot into the vRanger restore environment using DHCP.

To boot into the restore environment using DHCP

1   Using the boot media created in the previous steps, boot the target server into the vRanger restore environment.

2   When the boot is complete, you are presented with a list of available network devices and their connection status. If a DHCP server is available, the restore environment obtains its IP configuration automatically, and prompt with the IP information. Record this information to use during the restore job creation process.

3   If no DHCP server is detected, each ethernet device shows as "Disconnected". Highlight the preferred network device. If only one device is present, it is automatically highlighted. Press **Enter** to select the highlighted network interface. You are then prompted to select either **Use DHCP**, which attempts to obtain an IP address again, or **Configure Static IP**.

4   **[Optional]** If the vRestore environment does not contain the correct drivers for your hardware, you may enter a Linux command prompt by pressing **F1** and add them manually.

> ⓘ | **IMPORTANT:** Do not reboot the target server or the IP configuration is lost.

Booting into the restore environment using static IP

Complete the steps in the following procedure to boot into the vRanger restore environment using a static IP address.

To boot into the restore environment using static IP

1   Using the boot media created in the previous steps, boot the target server into the vRanger restore environment.

2   After the boot is complete, you are presented with a list of available network devices and their connection status.

- If no DHCP server is detected, each ethernet device shows as "Disconnected."

3   Highlight the preferred network device. If only one device is present, it is automatically highlighted. Press **Enter** to select the highlighted network interface. You are then prompted to select either **Use DHCP**, which attempts to obtain an IP address again, or **Configure Static IP**.

4   To configure a static IP address, select **Configure Static IP**. The **Setting Static IP** dialog appears.

5   Enter the IP address, subnet mask, and gateway information, and then select **Submit**. The target server is configured with the selected IP information.

6   [Optional] If the vRestore environment does not contain the correct drivers for your hardware, you may enter a Linux command prompt by pressing **F1** and add them manually.

> ⓘ | **IMPORTANT:** Do not reboot the target server or the IP configuration is lost.

Selecting a savepoint

The savepoint is from where the data for the physical machine is restored.

To select a savepoint

1   In the main vRanger interface, click **My Repositories**.

2   Select the repository in which the savepoints for the server you want to restore are stored.

3   In the **Working Repository** pane, select a savepoint to restore.

4   Do one of the following:

- On the toolbar, click **Restore**.
- Right-click the savepoint, and then and click **Restore**.

The vRanger Restore Wizard launches.

Naming the job

After the Restore Wizard - Advanced launches, you have the opportunity to name the restore job so that you can distinguish it from other jobs.

To name the job

1   On the Mode Selection page of the Restore Wizard, enter the **Job Name** and **Job Description** in the appropriate text boxes.

2   Click **Next**.

> ⓘ | **NOTE:** For the procedures on restoring from a manifest file, see Process of Performing a Full Restore from Manifest for Virtual Machines.

3   Populate the **Job Name** and **Job Description** text boxes.

4   Click **Next**.

The Disk Mapping page displays.

Mapping disks

The Disk Mapping page of the wizard lets you map the protected disks in the savepoint to the existing physical disks on the target server.

To map disks

1   On the Disk Mapping page, in the Restore Machine IP Address text box, enter the IP for the target server (see Booting the target server).

2   Click **Map Disks**.

The page lists the disks in the chosen savepoint.

3   For each disk you want to recover, select the appropriate target disk in the **Restore to Disk** drop-down menu.

- For disks that you do not want to recover, select **Skip this Disk**.

4 Click **Next**.

The Recurrence Schedule Selection page displays.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

To set the recurrence schedule

1 On the Recurrence Schedule Selection page, select from the options described in the following table.

This will be an On Demand job...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:

- **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
- **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
- **Monthly** – The monthly option offers the following configurations:
  - **Day *[x]* of every*[y]* month:**
    - *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
    - *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
  - **The *[f] [d]* of every *[y]* month(s):**
    - *f* can be either: first, second, third, fourth or last.
    - *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
- **Yearly** – The yearly option offers the following configurations:
  - **Every *[m][x]*:**
    - *m* is any month of the year. This determines the month of the year in which the report occurs.
    - *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
  - **The *[f][d]* of *[m]*:**
    - *f* can be either: first, second, third, fourth, or last.
    - *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
      m is any month of the year. This determines the month of the year in which the report occurs.
- **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between jobs.

NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date:

- **No end date** – The job recurs until it is stopped.
- **End after** – Select the number of occurrences.
- **End by** – Select the date by which to end the recurrence.

2   Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

    •   To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the restore job

The final page in the Restore Wizard displays a summary of the selections you made for the job.

To start the restore job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

    ⓘ   **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

    ⓘ   **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

    a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

    b   In the View toolbar, click **Run**.

    c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

About file-level restore - Windows

You can restore a file from a savepoint by accessing the File-Level Restore (FLR) command in the My Repositories view. FLR is accessible regardless of how the savepoints are sorted. You can right-click the savepoint in the Working Repository pane to select the command or you can click to highlight the savepoint and then click the FLR icon in the toolbar.

There is no need to copy or recreate the contents of the disk. The data remains compressed. When you make an FLR request, only the relevant files are moved from the repository to the designated computer.

A key component of the vRanger file-level recovery is the vRanger catalog functionality. While you can use FLR without cataloging enabled, it is much easier to find a preferred file by using a catalog search than through manual browsing of savepoints.

Limitations and requirements

There are several limitations and requirements that apply to file-level restores.

ⓘ   **NOTE:** These are summarized briefly below, see the *Dell vRanger System Requirements Guide* for complete information.

•   **Disk types** – VMs with dynamic disks are not supported for FLR.

Linux FLR limitations are summarized in Linux FLR limitations.

Performing a file-level restore using Catalog Search

Cataloging lets you search and browse savepoints for files inside the backup archives. Cataloging must be enabled at the start of the backup job to function, and also enabled globally in the Configuration Options dialog.

Catalog Search has the following parameters:

- When cataloging a full backup, all files in the system are stored in the database.

- When cataloging an incremental backup, only the files that were modified since the last backup are stored in the database.

- When cataloging a differential backup, only the files that were modified since the last full backup are stored in the database.

  > ⓘ | **NOTE:** When it comes to reporting the status of a backup task, cataloging is not considered a critical function. The cataloging portion of the task can fail, and the task is still considered successful.

To perform a file-level restore using Catalog Search

1 On the main vRanger interface, navigate to the My Repositories view, and then select the repository in which you want to search.

2 Do one of the following to launch the Catalog Search function:

   - Click the **Catalog Search Icon** 🔲.

   - On the Menu Bar, click the **Tools menu**, and then click **Catalog Search**.

   The Catalog Search & Browse dialog displays.

   Entering a search string (in the Catalog Search Criteria box) searches all repositories and savepoints for the string. To further refine the scope of the search, click **Advanced**.

3 The **Advanced** option lets you limit the search to a repository or VM (savepoint). Select the repository or VM you want to search, enter the search string, and click **Search**.

4 Highlight the preferred savepoint, and click **FLR** for File Level Restore, or **Restore** for the Full Restore.

   > ⓘ | **NOTE:** Catalog searching supports the traditional wildcard character (*) in any position. The search string can be as short or as long as you prefer; however, the shorter the string, the longer the search takes.SQL Server Express is limited to one CPU and 1 GB of RAM; an extremely short search string (for instance, "dot") could result in very slow searching, and SQL Server Express could run our of memory.To minimize performance issues during Catalog Search, make the search string as specific as possible.

Performing a manual file-level restore

The dialog box for this method of restoring files is made up of the following three panes:

- **Savepoint Pane** – In this area, all of the files or folders connected to a savepoint are nested. When you click the file or folder that you want to restore, the path is listed above it.

- **Destination Pane** – This pane includes a tree structure of all possible destinations for the restored file or folder you select.

- **Restore Status Pane** – This area of the dialog contains data after the restore process starts. It includes a column that indicates the percentage of the job that has completed. The dialog also features a *Stop* link that, when clicked, ends the restore process. When a job is stopped or it has completed, a *View in Explorer* link displays.

To perform a manual file-level restore

1 In the main vRanger interface, click **My Repositories**.

2 In the **Working Repository Pane**, right-click the savepoint for which you want an FLR and select **File Level Restore**.

3 In the **Select the Files/Folders to Restore** pane, expand the archive to view the file structure.

4 Select the file or files to recover.

5 In the **Destination** pane, select the destination to which the selected files should be recovered.

   The **Destination** pane shows the Windows Explorer view for the machine on which vRanger is installed.

⚠ | **WARNING: vRanger does not check for previous versions of the files you are restoring. Files restored using FLR overwrite any previous files (of the same name) in the restore destination.**

6 Click **Restore**.

7 In the **Restore Status** pane, view the status of the restore activity and the value in the Percent Finished column. Status should be Finished and Percent Finished should be 100%.

8 To view the restored file or folder, click the *View in Explorer* link.

About performing a Linux file-level restore

File-level restores (FLRs) from Linux server savepoints require that a vRanger virtual appliance is deployed and configured for Linux FLR. See the *Dell vRanger Pro Installation/Upgrade Guide* for instructions on deploying and configuring the vRanger virtual appliance.

After you configure the vRanger virtual appliance, the FLR process for a Linux VM is the same as the manual FLR process (meaning the process not using a catalog search) for a Windows VM. For more information, see Performing a manual file-level restore.

Linux FLR limitations

There are several limitations and requirements that apply to file-level restores from Linux VMs.

ⓘ | **NOTE:** These are summarized briefly below. See the *Dell vRanger System Requirements Guide* for more information.

- **Linux** – File-level recovery from Linux servers requires the use of the vRanger virtual appliance. See About performing a Linux file-level restore for more information.

- **Linux** – Not all Linux file systems are supported for file-level recovery. See the *Dell vRanger Systems Requirement Guide* for a list of supported Linux platforms and file systems.

- **Permissions** – vRanger requires you to recover the Linux files to an intermediate Windows machine. When you recover Linux files to a Windows machine, you lose the file permissions.

- **Source VM Configuration** – The source VM properties needs to show the OS type as Linux. If this is not configured properly, vRanger does not identify the savepoint as a Linux VM.

Restoring from manifest

vRanger offers the ability to restore a savepoint after it has been moved from a repository. This allows for greater flexibility when restoring a file or server after it has been moved to long-term tape backups. You may also use the restore from manifest feature to restore a server after transporting the savepoint to a secondary location.

You may perform three different operations when restoring from a manifest file:

- Process of performing a full restore from manifest for virtual machines
- Process of performing a full restore from manifest for a physical machine
- Performing a file-level restore from manifest

Process of performing a full restore from manifest for virtual machines

Each savepoint contains a manifest file that can be used to restore the savepoint without access to the vRanger database. If the vRanger machine is lost, you may install a fresh installation of vRanger, and restore savepoints using this manifest file.

> (i) | NOTE: The procedures in this section apply to restoring from a virtual machine savepoint manifest. For information on restoring a physical server from a savepoint manifest, see Process of Performing a Full Restore from Manifest for a Physical Machine.

To perform a full restore from manifest for virtual machines, complete the following tasks:

1. Creating a restore job
2. Selecting a savepoint from the manifest
3. Selecting a destination
4. Selecting a network
5. Selecting a transport method for VMware machines
6. Configuring restore options
7. Setting the recurrence schedule
8. Configuring email notifications
9. Starting the restore job

Creating a restore job

Creating a restore job launches the Restore Wizard, which guides you through the steps of completing the job.

To create a restore job

1. On the vRanger main interface, click **My Repositories**.
2. On the toolbar, click **Restore from Manifest**, and then click one of the following options:
   - Restore from VMware Savepoint Manifest File
   - Restore from Hyper-V Savepoint Manifest File

   The Restore Wizard launches.
3. On the Mode Selection page, populate the **Job Name** and **Job Description** text boxes.
4. Click **Next**.

   The Savepoint Manifest Selection page appears.

Selecting a savepoint from the manifest

Use the Savepoint Manifest Selection page to browse to the location of the manifest on your network, and configure the appropriate authentication credentials for the file.

To select a savepoint from the manifest

1   On the Savepoint Manifest Selection page, do one of the following:

   •   Enter the path to the savepoint manifest file.

   •   Click **Browse**, and then browse to the location of the savepoint manifest.

2   Enter the credentials for the account with access to the savepoint manifest.

3   [Optional] If the savepoint is encrypted, enter the encryption password.

4   Click **Add**.

5   Repeat steps 1-4 as needed for multiple savepoints.

6   Click **Next**.

   The Destination Selection page appears.

Selecting a destination

The Destination Selection page lets you configure the host, drives, and datastores to which the savepoints are restored.

To select a destination

1   To select options to be applied to all VMs, do the following:

   •   Select the host from the **Master Host** drop-down list.

   •   Confirm that the path that appears in the Virtual Machine Location text box is correct.

      To change this path, use the browse (...) button to locate and a select a new location.

   •   Click the Master Disk Storage Datastore browse button to locate and select the datastore to be used for all VMs.

   •   Click **Assign**.

2   To select options for individual VMs, do the following:

   •   Select the disks you want to restore.

      Disks not selected appear in italics.

   •   To select a different host for each VMDK, click the drop-down list in the Restore To Host column for each VM, and then select the host to which the VM should be restored.

   •   To select a different datastore for each VM, click the button next to the Primary Datastore column entry, and then browse for and select a datastore location.

3   To rename to VM you are restoring, enter the new name under the Restore As column.

4   Click **Next**.

Selecting a network

The Network Selection page lets you configure the virtual network to which each vNIC should be connected upon restore. You can also configure a Master Network, which uses the selected network for all vNICs.

To select a network

1   To use the same network for all VMs, do the following:

   a   Select a host from the **Master Network** drop-down list to which all VMs should be restored.

   b   Click **Assign**.

2 To use different networks for individual VMs, select a virtual network from the connected network adapters included in the **Restore To Host** drop-down list for each VM to which the vNIC should be connected.

3 Click **Next**.

Selecting a transport method for VMware machines

In vRanger, the transport determines how backup data is sent, and where the backup processing activity occurs (for more information, see VMware backup transport options). When configuring a backup job for a VMware VM, the Backup Wizard, by default, uses the Automatic Transport Selection option to automatically select the best transport method at run-time, but you can also configure the transport manually by using the Using the Custom Transport selection option.

ⓘ | **NOTE:** The Transport Selection page is not available when backing up a physical or Hyper-V virtual server. When backing up physical servers and Hyper-V virtual machines, vRanger sends the backup data over the network to the repository. Because this is the only method available for transporting data for physical and Hyper-V servers, the Transport Selection page of the Backup Wizard does not display.

To select a transport method for VMware machines, complete one of the following procedures:

- Using the Automatic Transport selection
- Using the Custom Transport selection

Using the Automatic Transport selection

The Backup Wizard includes an Automatic Transport Selection option, which lets vRanger select the best available method for your configuration. When determining the best transport, vRanger uses two key criteria:

- Whether virtual appliances (VAs) are deployed.
- Where vRanger is installed.

The logic used to select a transport differs slightly based on whether vRanger is installed on a VM or on a physical server. For more information about the order in which vRanger checks for transport availability, see Transport selection order - virtual machine backup.

To use the Automatic Transport selection

1 On the Transport Selection page, select **Automatic transport selection**.

The selected transport method displays at the bottom of the page.

ⓘ | **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method may not always be what is shown on the page during the job setup.

vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

- If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

For more information about configuring the VA, see Understanding the vRanger virtual appliance.

2 Click **Next**.

Using the Custom Transport selection

The Custom Transport Selection option lets you specify both where the actual backup processing activity occurs) and which transport to use. For more information, see VMware backup transport options.

To use the Custom Transport selection

1 On the Transport Selection page, select **Custom transport selection**.

This enables the custom settings.

2   Select where you want the job to occur from the following options:

- On this vRanger machine.

- On a vRanger virtual appliance.

vRanger verifies whether any virtual appliances are configured on the source host. If the host belongs to a cluster, and no host VA is found, vRanger searches for shared VAs on the cluster as well.

- If no VA is detected, click **Configure Virtual Appliance** to add a VA to vRanger.

  For more information about configuring the VA, see Understanding the vRanger virtual appliance.

  > ⓘ | **NOTE:** Using the vRanger virtual appliance is recommended for scalability and performance.

3   Select which transports the job should attempt from the following options:

- Advanced [SAN / HotAdd]

- LAN

For more information, see VMware backup transport options.

4   [Optional] If you want your LAN backups to be encrypted, select **Use SSL encryption when transferring files over the LAN connection**.

⚠ | WARNING: **This option causes reduced performance.**

5   Click **Next**.

Configuring restore options

The Options Selection page displays the restore options available depending on the type of VM you want to restore.

To configure restore options

1   On the Options Selection page, select from the following restore options:

- **Force overwrite** – This overwrites any existing disk or config files of the same name as those in the destination host.

- **Force power on** – This turns on the VM after the restore.

- **Create cluster role for restored virtual machine** – This option is for Hyper-V VMs only. It attempts to add the VM as a cluster resource of the Hyper-V host parent cluster.

  > ⓘ | **NOTE:** The cluster option is available only when a Hyper-V cluster has been added to the vRanger inventory and includes the host that is selected as the restore destination.

2   Click **Next**.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the restore job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

To set the recurrence schedule

1   On the Recurrence Schedule Selection page, select the options as described in the following table.

This will be an On Demand job...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:
	
	• **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
	
	• **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
	
	• **Monthly** – The monthly option offers the following configurations:
	
	    • **Day *[x]* of every *[y]* month:**
	        • *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
	        • *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
	
	    • **The *[f] [d]* of every *[y]* month(s):**
	        • *f* can be either: first, second, third, fourth or last.
	        • *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	        y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
	
	• **Yearly** – The yearly option offers the following configurations:
	
	    • **Every *[m][x]*:**
	        • *m* is any month of the year. This determines the month of the year in which the report occurs.
	        • *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
	
	    • **The *[f][d]* of *[m]*:**
	        • *f* can be either: first, second, third, fourth, or last.
	        • *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	        m is any month of the year. This determines the month of the year in which the report occurs.
	
	• **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between jobs.
	
	NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date: • **No end date** – The job recurs until it is stopped. • **End after** – Select the number of occurrences. • **End by** – Select the date by which to end the recurrence.

2   Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

    - To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the restore job

The final page in the Restore Wizard displays a summary of the selections you made for the job.

To start the restore job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

    ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

    ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

    a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

    b   In the View toolbar, click **Run**.

    c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

Process of performing a full restore from manifest for a physical machine

vRanger offers the ability to restore a savepoint after it has been moved from a repository. This allows for greater flexibility when restoring a file or server after it has been moved to long-term tape backups. You may also use the restore from manifest feature to restore a server after transporting the savepoint to a secondary location.

Each savepoint contains a manifest file that can be used to restore the savepoint without access to the vRanger database. If the vRanger machine is lost, you may install a fresh installation of vRanger, and restore savepoints using this manifest file.

ⓘ | **NOTE:** The procedures in this section apply to restoring from a physical machine savepoint manifest. For information on restoring a virtual machine from a savepoint manifest, see Process of Performing a Full Restore from Manifest for Virtual Machines.

To perform a full restore from a manifest for a physical machine, complete the following tasks:

1　Creating the boot media

2　Booting the target server

3　Creating a restore job

4　Selecting a savepoint from the manifest

5　Mapping disks

6　Setting the recurrence schedule

7　Configuring email notifications

8　Starting the restore job

Creating the boot media

To restore a savepoint to a physical machine, you must first boot the target server into the vRanger restore environment using the vRanger Restore images (ISO for CD, IMG for a bootable USB drive).

When vRanger is installed, two restore image files (*vzRestore-0.1.2.iso* and *vzRestore-0.1.2.img*) are extracted to the "Service" sub-directory of the vRanger installation folder. By default this is on of the following:

- On x86 machines – *C:\Program Files (x86)\Dell\vRanger\Service*
- On x64 machines – *C:\Program Files\Dell\vRanger\Service*

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To create the boot media, complete one of the following procedures:

- Creating a bootable CD
- Creating a bootable USB drive

Creating a bootable CD

Complete the steps in the following procedure to create a bootable CD.

> ⓘ IMPORTANT: The procedures for creating a bootable CD differ widely depending on the operating system and software used. The procedures below show how to use the native Windows 7 utility to create a bootable CD. If you have a different image burning application, follow the instructions for that application.

> ⓘ NOTE: This procedure requires a blank CD/DVD and a writable CD/DVD drive.

To create a bootable CD

1　Insert the blank CD/DVD into the writable drive.

2　In the vRanger installation directory, find the file *vzRestore-1.2.6.iso*. If necessary, copy this file to the machine with the writable CD/DVD drive.

3　Right-click on the ISO file, select **Open with**, and then click **Windows Disk Image Burner.**

4　From the disk burner drop-down list, select the drive containing the blank disk.

5　Click **Burn**.

Creating a bootable USB drive

Complete the steps in the following procedure to create a bootable USB drive.

> **NOTE:** There are no native Windows utilities for creating a bootable USB drive from an IMG file. The procedure below was created using a freeware application - Image Writer for Windows. Follow the instructions for your application.

To create a bootable USB drive

1 Insert a blank USB drive into your USB port.

2 In the vRanger installation directory, find the file *vzRestore-1.2.6.img*. If necessary, copy this file to the machine with the blank USB drive.

3 In the Image Writer for Windows UI, browse to the *vzRestore-1.2.6.img* file.

4 In the **Device** drop-down list, ensure that the correct USB drive is selected.

5 Click **Write**.

6 Click **Exit**.

Booting the target server

In order for vRanger to detect the target server, the target server must be booted into the vRanger restore environment and have a valid IP address. The vRanger Restore image is configured to use Dynamic Host Configuration Protocol (DHCP) by default. If you are using static IP, some additional steps are required.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To boot the target server, complete one of the following procedures:

- Booting into the restore environment using DHCP

- Booting into the restore environment using static IP

Booting into the restore environment using DHCP

Complete the steps in the following procedure to boot into the vRanger restore environment using DHCP.

To boot into the restore environment using DHCP

1 Using the boot media created in the previous steps, boot the target server into the vRanger restore environment.

2 When the boot is complete, you are presented with a list of available network devices and their connection status. If a DHCP server is available, the restore environment obtains its IP configuration automatically, and prompt with the IP information. Record this information to use during the restore job creation process.

3 If no DHCP server is detected, each ethernet device shows as "Disconnected". Highlight the preferred network device. If only one device is present, it is automatically highlighted. Press **Enter** to select the highlighted network interface. You are then prompted to select either **Use DHCP**, which attempts to obtain an IP address again, or **Configure Static IP**.

4 **[Optional]** If the vRestore environment does not contain the correct drivers for your hardware, you may enter a Linux command prompt by pressing **F1** and add them manually.

> **IMPORTANT:** Do not reboot the target server or the IP configuration are lost.

Booting into the restore environment using static IP

Complete the steps in the following procedure to boot into the vRanger restore environment using a static IP address.

To boot into the restore environment using static IP

1   Using the boot media created in the previous steps, boot the target server into the vRanger restore environment.

2   When the boot is complete, you are presented with a list of available network devices and their connection status.

    - If no DHCP server is detected, each ethernet device shows as "Disconnected."

3   Highlight the preferred network device. If only one device is present, it is automatically highlighted. Press **Enter** to select the highlighted network interface. You are then prompted to select either **Use DHCP**, which attempts to obtain an IP address again, or **Configure Static IP**.

4   To configure a static IP address, select **Configure Static IP**. The **Setting Static IP** dialog appears.

5   Enter the IP address, subnet mask, and gateway information. Select **Submit**. The target server is configured with the selected IP information.

6   **[Optional]** If the vRestore environment does not contain the correct drivers for your hardware, you may enter a Linux command prompt by pressing **F1** and add them manually.

    ⓘ | **IMPORTANT:** Do not reboot the target server or the IP configuration is lost.

Creating a restore job

Creating a restore job launches the Restore Wizard, which guides you through the steps of completing the job.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To create a restore job

1   On the vRanger main interface, click **My Repositories**.

2   On the toolbar, click **Restore from Manifest**, and then click Restore from **Physical Savepoint Manifest File**.

    The Restore Wizard launches.

3   On the Mode Selection page, populate the **Job Name** and **Job Description** text boxes.

4   Click **Next**.

    The Savepoint Manifest Selection page appears.

Selecting a savepoint from the manifest

Use the Savepoint Manifest Selection page to browse to the location of the manifest on your network, and configure the appropriate authentication credentials for the file.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To select a savepoint from the manifest

1   On the Savepoint Manifest Selection page, do one of the following:

    - Enter the path to the savepoint manifest file.
    - Click **Browse**, and then browse to the location of the savepoint manifest.

2   Enter the credentials for the account with access to the savepoint manifest.

3   [Optional] If the savepoint is encrypted, enter the encryption password.

4   Click **Add**.

5   Repeat steps 1-4 as needed for multiple savepoints.

6   Click **Next**.

The Destination Selection page appears.

Mapping disks

The Disk Mapping page of the wizard lets you map the protected disks in the savepoint to the existing physical disks on the target server.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To map disks

1 On the Disk Mapping page, in the Restore Machine IP Address text box, enter the IP for the target server (see Booting the target server).

2 Click **Map Disks**.

The page lists the disks in the chosen savepoint.

3 For each disk you want to recover, select the appropriate target disk in the **Restore to Disk** drop-down menu.

   • For disks that you do not want to recover, select **Skip this Disk**.

4 Click **Next**.

The Recurrence Schedule Selection page displays.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the restore job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To set the recurrence schedule

1 On the Recurrence Schedule Selection page, select the options as described in the following table.

This will be an On Demand job...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:
	
	• **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
	
	• **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
	
	• **Monthly** – The monthly option offers the following configurations:
	
	    • **Day *[x]* of every*[y]* month:**
	
	        • *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
	
	        • *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
	
	    • **The *[f] [d]* of every *[y]* month(s):**
	
	        • *f* can be either: first, second, third, fourth or last.
	
	        • *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	         y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
	
	• **Yearly** – The yearly option offers the following configurations:
	
	    • **Every *[m][x]*:**
	
	        • *m* is any month of the year. This determines the month of the year in which the report occurs.
	
	        • *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
	
	    • **The *[f][d]* of *[m]*:**
	
	        • *f* can be either: first, second, third, fourth, or last.
	
	        • *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	         m is any month of the year. This determines the month of the year in which the report occurs.
	
	• **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between jobs.
	
	NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the starting date and ending date: • **No end date** – The job recurs until it is stopped. • **End after** – Select the number of occurrences. • **End by** – Select the date by which to end the recurrence.

2 Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

      • To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the restore job

The final page in the Restore Wizard displays a summary of the selections you made for the job.

This task is a procedure in Process of performing a full restore from manifest for a physical machine.

To start the restore job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

      ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

      ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

      a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

      b   In the View toolbar, click **Run**.

      c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

Performing a file-level restore from manifest

A file-level restore from manifest lets you move the files that make up a backup to your preferred location restore a virtual machine whose backup files are not in a repository. This lets the user move the files that make up a backup to wherever they want and still perform a restore.

To perform a file-level restore from manifest

1   Click **My Repositories**.

2   In the My Repositories toolbar, click the **FLR from Manifest icon** 🗂. The File Level Restore from Manifest dialog appears.

3   Enter the user name, password and path for the Manifest file of the backup that contains the files you want to restore.

4   Enter the Encryption Password, if the backup was encrypted.

5   Click **OK**.

6   In the Select the Files/Folders to Restore pane, expand the archive to view the file structure, and then select the file or files to recover.

7   In the Destination pane, select the destination to which the selected files should be recovered.

The Destination pane shows the Windows Explorer view for the machine on which vRanger is installed.

⚠ | **WARNING: vRanger does not check for previous versions of the files you are restoring. Files restored using FLR overwrites any previous files (of the same name) in the restore destination.**

8   Click **Restore**.

9   In the Restore Status pane, view the status of the restore activity and the value in the Percent Finished column. Status should be Finished and Percent Finished should be 100%.

10 To view the restored file or folder, click **View in Explorer**.

Replicate

- Understanding replication
- Replication modes
- Configuring transport failover options
- Creating a replication job
- About failover
- Performing failback

Understanding replication

vRanger includes integrated replication based on the proven technology of vReplicator, enabling replication of VMware virtual machines both on site and at remote locations for flexible and efficient disaster recovery preparedness. Combine backup and replication jobs to ensure you meet your organization's recovery time and recovery point objectives. Perform fast VM failover and recovery at all of your sites, no matter where they are located.

> ⓘ | **NOTE:** Replication is not available for Hyper-V virtual machines.

A VM is made up of a set of files, which means that replicating a VM is, in essence, replicating the set of files that make up the VM, with changes to these files that reflect user specified settings for the source VM.

The set of files replicated by vReplicator is listed below:

Table 7. Files replicated

File extension	Description
.vmx	The VM config file, one per VM
.vmxf	The extended VM config file, one per VM
.nvram	The VM BIO file, one per VM
.vmdk	The VM hard disk file, one per hard disk or snapshot
-flat.vmdk	The VM hard disk data file, one per hard disk or snapshot
-delta.vmdk	The VM snapshot data file, one per snapshot
-ctk.vmdk	The VM hard disk change tracking file when CB is enabled on the disk

How replication works

During the replication process, the configuration files are created and modified on the target server by way of the VMware API.

A set of working files is also created and used during the replication process. These files and their purposes are listed below:

Table 8. Working files

File extension	Description
.vzmap	Records data block offset and hash of files on the target VM. A vzmap file is created for each of the files replicated at the end of the replication. The vzmap file is used by the next replication pass to detect any data changes since the previous pass. It stays on the target VM as long as the job is still configured to run. While relatively small, the size of the vzmap file is directly proportional to the size of the VMDK it is based on. During replication, the vzmap file is stored on the target VA.
vzundo-script	A script created by the replication process to roll back changes on the target VM in case of a replication failure. This file created on the target VA at the beginning of the replication process and removed at the end.
.vzundo	This is a temporary file that records original data of changed blocks since the replication started. One for each VMDK replicated. In case of a replication failure such as network failure, the vzundo-script can be executed to restore files to their original state. These files are created on the virtual appliance by the replication process and removed when the job is completed. NOTE: The vzundo file is as large as the amount of changed data replicated during a given pass. For example, if a replication pass sends 20 GB of changed data to the target VA, the vzundo file is 20 GB.
.vmdk-abbt.vztemp	Active block filter file. One for each hard disk data file when ABT is enabled. It records active data block offsets for source VM disks. This file is used against the.vzmap file to figure out data blocks that need to be streamed to the target. It is created at the beginning of the replication process and removed when disk replication is completed.
.vmdk-abfg.vztemp	Change block filter file. One for each hard disk when CBT is enabled. It records changed block offsets for source VM disks. This file is only generated when CBT and ABT are both enabled. It is later be combined with the .vmdk-abbt.vztemp file into -flat-map.vztemp and removed.
-flat-map.vztemp	Disk data filter file. One for each hard disk when one of two situations are true: CB is enabled, or both AB and CB are enabled. It contains active/changed data block offsets that need to be compared to the .vzmap file at the target to figure out data blocks that need to be streamed to the target. It is created right before file replication starts and removed when file replication is completed.
.vzcid	Records target VM disk CIDs at the end of the replication pass.

Replication with the virtual appliance

vRanger supports ESXi replication by way of the vRanger Virtual Appliance, which leverages VMware's HotAdd disk transport mechanism. After the virtual appliances are configured and deployed, the use of the virtual appliance is automatic and transparent. Some key points about replicating with the virtual appliance are listed below:

- See the *Dell vRanger Installation/Upgrade Guide* for instructions on deploying and configuring the virtual appliances.

- You need to deploy, at minimum, one virtual appliance on every ESXi server that you plan to use as a replication source or target. For hosts in a cluster, you may deploy the VA to just one host in the cluster; the VA is shared among the cluster's other hosts.

  Virtual appliances must be used in pairs. For example, if you are replicating to a target host or cluster that is using a virtual appliance, you must use a virtual appliance on the source host/cluster as well.

- Due to VMware limitations on how snapshots are handled during HotAdd operations, hardware changes to the source VM occurring after a job has been configured are not replicated to the target VM. See Replication modes below.

Communication between the virtual appliances occurs through an SSH tunnel using AES-256 encryption. See the **Encryption** description in Major feature list for more information.

Additional replication requirements

The following limitations and requirements apply to replication:

- vRanger replication will not operate with servers that are behind a NAT firewall. In order to properly replicate through an NAT firewall, you must have an IP tunnel in place between two NAT'ed subnets. Contact your ISP provider to see if this option is available to you.

- The virtual machine hardware cannot be changed during replication. For this reason, the VM must be at a hardware version level that is compatible with both the source and target servers. The ESXi version of the source and target hosts does not matter, as long as the VM hardware is supported on both ESXi versions. For more information on VM hardware versions, consult the VMware documentation topic "Virtual Machine Compatibility".

- A continuous connection between source and target sites is required when replication is taking place.

- Excessive network packet loss could result in replication failure.Replication will work with links having average packet loss of less than 2%. Replication is not designed to work in replication environments where packet loss can exceed commercially accepted limits.

- Networks having 99% Uptime/Availability will generally provide for good Replication performance.

Replication with user snapshots

When replicating a VM that contains user snapshots, vRanger replicates all of the snapshots in the chain from the current snapshot to the base disk. At the target side (the VM to which changes are replicated), the snapshots are merged into a single disk.

Snapshots not in the chain of the current snapshot are not replicated. In the image below, ss2 is the current snapshot. Only ss1 and ss2 are replicated. Snapshot 3 (ss3) is a lateral snapshot to ss2, while ss4 is a child to ss2.
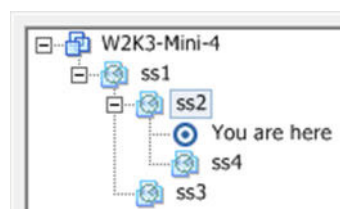
> ⓘ | **NOTE:** Replication with user snapshots applies only to replication using the virtual appliance. Service Console replication tasks cannot process user snapshots. If you need to replicate a VM with user snapshots, use a virtual appliance.

Capturing hardware changes with replication

It is not possible to replicate hardware changes that occur on the source VM after the job has been configured. This is due to a limitation in how snapshots are processed by HotAdd.

If you make hardware changes on the source VM, you need to manually configure the target VM in the same way b*efore* the next replication pass. You may also need to edit the replication job to include the new hardware.

Active Block Mapping (ABM)

ABM filters deleted data blocks so that only active blocks are scanned and streamed to the target. White space detection eliminates the need to compress, stream, and write zero blocks during the replication process. vRanger offers the following ABM settings options:

- Do not include the Windows page file

- Do not include the Windows hibernation file

Replication modes

vRanger offers Replication with Changed Block Tracking or standard Differential replication. Each of these replication modes has the option of Active Block Mapping (ABM). VM replication in general starts with replicating the source VM to the target host. Changes are applied to the target VM at user designated intervals to keep the target in sync with the source. Thus the key difference between the replication modes is how vRanger identifies VM changes to replicate. Refer to the sections below for more information:

- Replication with Changed Block Tracking
- Differential replication

Replication with Changed Block Tracking

Replication with Changed Block Tracking (CBT) is the recommended method of replication whenever possible.

CBT tracks the disk block changes made by the source VM. When enabled on the source host, CBT records the blocks that have changed since the last replication pass and transfers them to the target host without scanning the VMDK. CBT only identifies these disk sector changes when ESX version 4.1 or later is installed on both of the hosts being used for replication.

The following configurations are required for CBT to be available:

- The VM to which the disks belong must use Hardware Version 7 or later.
- CBT must be enabled for the VM you want to replicate.
- VM storage must not be located on an independent disk that is not affected by snapshots.

  > ⓘ **NOTE:** NFS is supported. RDM is supported in Virtual Compatibility mode, but not in Physical Compatibility mode. VMFS is also supported, regardless of whether it is backed by a SAN, iSCSI, or local disk.

For CBT to identify disk sectors in use with the special change ID, the following configurations are also required:

- The virtual disk must be located on a VMFS volume, backed by a SAN, iSCSI, or local disk.

  > ⓘ **NOTE:** Raw device map (RDM) partitions are not considered VMFS volumes.

- The VM must have zero (0) snapshots when CBT is enabled for a clean start.

  > ⓘ **NOTE:** CBT must be enabled for each VM that you want to replicate. vSphere supports CBT, and most VMs running in this environment can use it.

Differential replication

In differential replication, vRanger compares the data blocks on the source VM to a data map to identify blocks that have changed since the last replication pass. These changed blocks are sent to the target VM and committed to disk, and the data map is updated to reflect the current state of both VMs.

When using differential replication, consider the following:

- During the first replication pass, the entire VM image needs to be sent to the target VM. During subsequent replication passes, only changed data is sent. To avoid sending the entire VM image over the network, consider pre-seeding the target VM. For more information, refer to Pre-seeding replication jobs
- After the first replication pass, the source VM is scanned and the blocks are compared to the most recent data map. This scanning process can take an average of one minute per gigabyte of hard drive space; however, the actual speed depends on the underlying hardware being used. While only the changed data

is sent to the target VM, the entire hard drive must be scanned. When configuring differential replication jobs for large VMs, take this scanning time into account. To eliminate the need for scanning, use replication with Changed Block Tracking if at all possible.

Configuring transport failover options

Replication transport actions can occur either in the Service Console or through HotAdd (Lan-Free or Network). If one transport method is not available, vRanger attempts to perform the replication task using one of the other transport methods. You may configure the order in which vRanger attempts the transport options with the Transport Failover node of the Configuration Options page.

The recommended transport order for ESXi replication is ESXi/ESX 5 **VA-based HotAdd> VA-based Network**.

Pre-seeding replication jobs

Ranger replication is intended to replicate changes from a source VM to a remote target. It is often not practical to perform the first replication pass (which sends the full VM) to a remote site over a WAN link. You may use vRanger to "seed" a replication job locally to reduce the amount of data sent over the WAN.

To pre-seed replication jobs, complete the following tasks:

1   Backing up the source VM

2   Replicating the changes

Backing up the source VM

Complete the steps in the following procedure to back up the source VM.

To back up the source VM

1   If vRanger is not already installed, do so.

2   Select the VM you want to get to the DR site and create a Full backup job.

3   Backup the VM locally.

4   Copy the backup savepoint, and move it to the DR site. This is usually accomplished by shipping a USB drive or some other media containing the backup archive.

   Savepoints are named using a *VMname_date_time_UUID* format. Copy the entire directory (containing a .var file, and two .metadata files) for the savepoint in question.

5   Using the Restore from Manifest procedure documented in Process of performing a full restore from manifest for virtual machines, restore the backup to the target ESX Server at the DR site.

Replicating the changes

Complete the steps in the following procedure to replicate the changes made to the VM.

To replicate the changes

1   Launch vRanger and expand the host where the source VM is located. Select the VM.

2   Right-click and select **Replicate [VM name]**.

3   Configure a replication job with the preferred options.

4   When selecting the target host, select the host on the target side to which the source VM has been restored.

> **NOTE:** When selecting the target VM name, ensure that you are using the same name as the source VM. vRanger appends "_VzReplicate" to the VM name on the target side.

5   On the Select Target Host dialog, next to **Preseed VM UUID**, click the browse button and select the pre-seeded VM.

6   When configuring the primary datastore, ensure it is the same datastore as the restored VM and click **Finish**.

7   Run the replication job. Only the data changes pass over the WAN as opposed to the full VM.

Creating a replication job

vRanger uses a simple wizard to guide you through the process of creating a replication job. At a high level, the process includes the following steps. Each of the steps below is explained in more detail in the sections that follow.

To create a replication job, complete the following tasks:

1   Selecting the source VM

2   Selecting a target host

3   Selecting the replication mode

4   Selecting replication network mapping

5   Selecting hard disks

6   Selecting replication options

7   Setting the recurrence schedule

8   Configuring email notifications

9   Starting the replication job

Selecting the source VM

Complete the steps in the following procedure to select the source VM.

To select the source VM

1   Click **My Inventory**.

2   Right click the VM you want to Replicate, and select **Replicate [VM Name]**.

> **NOTE:** If there is no virtual appliance provided for the selected VM, the Virtual Appliance Not Found window notifies you that the replication job is going to fail without a virtual appliance. For more information, see Understanding the vRanger virtual appliance.

3   Enter the **Job Name** and **Description**, then click **Next**.

Selecting a target host

On the Select Target Host page, complete the following steps to select a target host.

To select a target host

1   Do one of the following:

   •   Choose a **Target Host** from the list, then click **Next**.

- If you have pre-seeded the VM on the target host, select the host on the target side to which the source VM has been restored. See Pre-seeding replication jobs for more information. When selecting the target VM name, ensure that you are using the same name as the source VM. vRanger appends "_VzReplicate" to the VM name on the target side.

2   [Optional] On the Select Target Host dialog, next to **Preseed VM UUID**, click the browse button and select the pre-seeded VM.

> ⓘ **IMPORTANT:** After a replication job has been saved, you may not change the target host. If you want to replicate to a different host, you may either vMotion the VM to the preferred host, or create a new replication job for the preferred target. If you vMotion the VM, the existing replication job should identify the change in host and adjust automatically.

Selecting the replication mode

On the Replication Mode page, you can select between differential replication and hybrid replication. For more information, see Replication modes.

To select the replication mode

1   On the Replication Mode page, under Replication Type, select **Differential Replication**.

> ⓘ **NOTE:** Hybrid replication is not supported when replicating to or from VMware ESXi. This is a deprecated feature for vRanger 7.2 and beyond.

2   Click **Next**.

Selecting replication network mapping

In the Replication Network Mapping page, you specify the network on the target host to which the VM nics should connect.

To select replication network mapping

1   Select the **Network Mapping** you want to use for this job.

2   Click **Next**.

Selecting hard disks

The Hard Disk selection page lets you include or exclude specific disks from replication.

To select hard disks

1   Select the Hard Disk(s) you want to replicate

2   You may select a **Primary Datastore,** which is used as the default value for all disks in this job, letting you configure only the exceptions.

3   In the **Target Datastore** column, verify that the Primary Datastore has been populated. If you want the disk to replicate to a different datastore, change it in this column.

4   Click **Next**.

Selecting replication options

The Replication Options page lets you configure additional operations performed during the replication job. You may also view the transport failover preferences, and modify them as needed.

To select replication options

1   Configure the replication options, using the table below for guidance.

Table 9. Guidelines

Option	Description
Check destination for free space.	This option prompts vRanger to check the destination location to ensure that there is enough space available for the replica VM.
Update notes with latest replication results.	If this flag is selected, vRanger updates the VM Notes in VC with the most recent replication task status. **NOTE:** vRanger maintains the existing VM notes in vCenter, appending the replication status to the end.
Enable guest quiescing.	This option enables the VMware Tools when creating a snapshot. Use this flag if you are replicating a database server. During this process, vRanger freezes or pauses writes on supported applications to provide a consistent image of VSS-enabled databases. See Physical backup transport options for more information.
Enable Active Block Mapping (ABM)	Checking this box enables Active Block Mapping, which lets vRanger skip data that recognized as removed on Windows NTFS (including junk data in recycle bin). **NOTE:** With this option enabled, junk data in recycle bin and deleted data on the VM are not replicated. The following ABM settings options are also available: • Do not include the Windows page file • Do not include the Windows hibernation file

2   Verify your Transport Failover preference. See Configuring transport failover options for more information. To change the Transport Failover order, click **Modify**. Otherwise, click **Next**.

Setting the recurrence schedule

The Recurrence Schedule Selection page lets you indicate whether the job should be done on demand (at this moment), or at a future time. To schedule a future occurrence, you must define the start time, recurrence pattern, and range of recurrence.

ⓘ **NOTE:** When using the Full Backup Days option, ensure that the selected recurrence schedule causes backup tasks to be run on the configured Full Backup Days.

To set the recurrence schedule

1   On the Recurrence Schedule Selection page, select from the options described in the following table.

ⓘ **NOTE:** To replicate more frequently than once per day, use the Interval reccurrance pattern.

This will be an On Demand job...	There is no need to define a schedule. The job runs on demand.
Start Time	The time of day at which to run the job.

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:
	
	• **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
	
	• **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
	
	• **Monthly** – The monthly option offers the following configurations:
	
	   • **Day** *[x]* **of every** *[y]* **month:**
	
	      • *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
	
	      • *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
	
	   • **The** *[f] [d]* **of every** *[y]* **month(s):**
	
	      • *f* can be either: first, second, third, fourth or last.
	
	      • *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	      y can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
	
	• **Yearly** – The yearly option offers the following configurations:
	
	   • **Every** *[m][x]:*
	
	      • *m* is any month of the year. This determines the month of the year in which the report occurs.
	
	      • *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
	
	   • **The** *[f][d]* **of** *[m]:*
	
	      • *f* can be either: first, second, third, fourth, or last.
	
	      • *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
	      m is any month of the year. This determines the month of the year in which the report occurs.
	
	• **Interval** – The Interval option offers the following configurations:
	
	   • **Day(s).** The default value is 0, which lets you replicate more frequently than daily.
	
	   • **Hour(s).** The default value is 1. To replicate more frequently, change to "0" and configure a "minutes" value.
	
	   • **Minute(s).** The default value is "0". The minimum interval is five minutes.
	
	NOTE: Configure a value sufficient to let the previous replication pass to continue.

Range of Recurrence	Select the starting date and ending date. • **No end date:** The job recurs until it is stopped. • **End after:** Select the number of occurrences. • **End by:** Select the date by which to end the recurrence.

2   Click **Next**.

Configuring email notifications

The Email Notification Selection page lets you select the email addresses to which notifications should be sent for this job. You can also request that a report be sent to the recipients identified after the job is complete.

To add or edit email addresses, see Setting up the SMTP server.

To configure email notifications

1   On the Email Notification Selection page, in the Selected column, select the address to which notifications should be sent.

    •   To select all email addresses, select the check box in the Selected column header.

2   Click **Next**.

Starting the replication job

The final page in the Replication Wizard displays a summary of the selections you made for the job.

ⓘ | **NOTE:** The most common configuration error is configuring a Transport Failover sequence that does not match your hardware. For example, if you have VA-HotAdd selected in the Transport Failover section, but one of your hosts does not have a VA configured.

To start the replication job

1   On the Summary page, review the data displayed in the tree view.

2   If you want to automatically begin the job immediately after completing the wizard, select **Run the Job after Finish is clicked**.

    ⓘ | **NOTE:** If this option is not selected, the job is saved and begins at the time determined during the Setting the Recurrence Schedule procedure.

3   Click **Finish**.

4   If you did not opt to automatically begin the job after clicking Finish, and this is an on-demand job, do the following:

    ⓘ | **NOTE:** If this is not an on-demand job, the job begins at the scheduled time.

    a   In the main vRanger interface, in the My Inventory view, select the server or host for the server you want to restore, and then select the job from the Working Inventory pane.

    b   In the View toolbar, click **Run**.

    c   In the Confirm Run Job dialog box, click **Yes**.

5   To check progress or confirm completion, click **My Jobs**, and then click **Recent Jobs**.

About failover

There are two kinds of replication Failover: Test Failover and Failover.

•   **Test Failover:** This function disables the replication job, snapshots the replica VM, and powers it on. You have the option of performing a test failover with or without networking enabled on the replica VM. There are two methods for performing failure:

    •   Performing Test Failover with networking disabled

    •   Performing Test Failover with networking enabled

- **Failover:** This simplifies the action of failover to a replica VM. You have the option to perform a final synchronization or to failover immediately. The vReplicator process gracefully and safely powers down your production VM, and boots your replica VM with a single click.

Performing Test Failover with networking disabled

The Test Failover operation can be performed from the Scheduled Jobs and On Demand Jobs nodes of the My Jobs view.

The default Test Failover operation is described below. To perform a test failover with networking enabled, refer to Performing Test Failover with networking enabled.

To perform test failure with networking disabled

1   Click the **Test Failover** icon. A confirmation message displays.

2   Click **Yes**.

     The test failover process begins, and a job is added to the Current Jobs node.

3   When testing is complete, select the Test Failover task in the Current Jobs node, and click **Resume**.

4   Click **Yes** when the confirmation prompt appears.

     The snapshots made are reverted.

     The network card will be "Connect at Power On". This prepares the replica VM for failover.

     The VM returns to it's dormant state.

Performing Test Failover with networking enabled

Performing a test failover with the network enabled has the potential to negatively affect your production environment. For this reason, this method of test failover is not recommended. The recommended method is Performing Test Failover with networking disabled.

ⓘ | **NOTE:** Due to this option's potential for negatively impacting your production environment, it is disabled by default.

To perform network-enabled test failover, complete the following tasks:

1   Enabling network-enabled Test Failover

2   Performing a network-enabled Test Failover

Enabling network-enabled Test Failover

This feature is enabled by way of the configuration file for the vRanger client - *Vizioncore.vRanger.Client.Shell.exe.config*. This file is located in the vRanger installation directory. By default this is:

```
C:\Program Files\Dell\vRanger
```

To enable network-enabled test failover

1   Open the *Vizioncore.vRanger.Client.Shell.exe.config* file with your text editor of choice.

2   Carefully make the required changes, using the information below as guidance.

     ⓘ | **NOTE:** Not all text editors include line numbers. If you do not see the line numbers referenced below in your text editor, this is not an indication of a problem.

3   Scroll to line 38 or search for **"Test Failover"**.

4 Change the value for **ShowTestFailoverDialog** from "false" to "true", ensuring the quotation marks remain.

5 Save the file and exit.

6 Restart the vRanger Service for the configuration change to take effect.

Performing a network-enabled Test Failover

After it is enabled, you are prompted to perform a test failover; you can do so with or without networking enabled during the test.

To perform a test failover

1 Click the **Test Failover** icon. The Test Failover dialog appears.

2 Select the preferred option, as described below:.

- **Continue with networking on the replica disabled** – When performing a test failover with networking disabled, the replica VM is powered on with the network adapter disconnected. This protects the production environment, but eliminates the ability of testing your replica VM as part of a larger environment.

- **Continue with networking on the replica enabled** – When performing a network enabled test failover, your replica VM is powered on with the network adaptor connected.

⚠ | **WARNING: This has the potential to impact your production environment - exercise caution when using this option.**

3 Click **OK**.

4 When testing is complete, select the Test Failover task in the Current Jobs node, and click **Resume**.

- The snapshots made are reverted.

- The network card is set to "Connect at power on". This prepares the replica VM for failover.

- The VM returns to its dormant state.

5 Click **Yes** when the confirmation prompt appears

Performing a failover

If the production site becomes unavailable, or you need to move the workload to the DR site, you have the option of performing a failover operation from the vRanger interface.

A failover operation consists of the following actions:

- Turning on the production VM.

- **[Optional]** Performing a final synchronization between sites. This ensures that no data is lost during a failover. If you do not want to synchronize changes (in the event of a virus, perhaps) or cannot (in the event of a hardware failure), data changed since the last replication pass is lost.

- Turn on the D/R VM.

- Reverse the replication job, with the D/R VM now acting as the source.

  ⓘ | **NOTE:** The first replication pass after failing over will send the entire image, not just the changed data.

To perform a failover

1 In the My Jobs view, select a replication job.

2 Click the **Failover** icon.

A confirmation message displays. Click **Yes**.

3   The Confirm Synchronization dialog appears.

- **Yes** – This shuts down the production VM and perform a final synchronization to the D/R site. This ensures that no data is lost, but also transfers any virus or corruption to the replica VM. View the replication progress in the Running Tasks section of the vRanger UI.

  ⓘ | **NOTE:** If you are using differential replication, this causes vRanger to perform a complete scan of the VM disks (as it does for every replication pass), but only changed data is sent to the replica VM.

- **No** – The production VM is turned off, but no changes are sent to the D/R replica.

4   After the Failover completes, the Destination of the original replication job is now the Source, and the original Source is now the Destination. In other words, the job is now set to replicate from the D/R site to the Production site. This change is not reflected in the vRanger UI. Jobs that have been failed over are indicated by a failover icon and tool-tip message:

5   If the production site is unavailable, or you do not want to replicate changes, disable the replication job.

Performing failback

When the production server has been returned to proper operation, you usually need to failback the replicated VM to the production site. A failback is simply a failover back to the production site.

A failback operation involves the following events:

- Turning off the D/R VM.

- [Optional] Performing a final synchronization between sites. This ensures that any data changed on the D/R VM is replicated back to the production site. If you do not synchronize changes, data changed since the last replication pass is lost.

- Turn on the production VM.

- Reverse the replication job (return to the jobs original configuration), with the D/R VM now acting as the target.

- If the job is enabled, replication continues, sending changes from the source production VM to the D/R target.

  ⓘ | **NOTE:** The first replication pass after failing back re-scans the entire image to identify changed blocks. Only changed data is sent.

To perform a failback

1   In the My Jobs view, select a replication job.

2   Click the **Failback** 🔲 icon.

A confirmation message displays.

3   Click **Yes**.

A **Confirm Synchronization** dialog appears.

4   Select one of the following:

- **Yes** – This ensures that any data changed on the D/R VM is replicated back to the production site. If you do not want to synchronize, data changed since the last replication pass is lost. View the replication progress in the Running Tasks section of the vRanger UI.

- **No** – The D/R VM turned off, but no changes are sent to the Production site.

After the Failback completes, the direction of replication reverses again, this time returning to the original configuration. The Failover icon is no longer shown.

Reports

- About vRanger reports
- My Reports view
- Generating a report using the Report Wizard

About vRanger reports

The My Reports view is a gateway to creating, editing, and distributing custom reports. There are several report types:

- Backup Job Summary Report
- Backup Task Report
- Replicate Job Summary Report
- Replicate Task Report
- Restore Job Summary Report
- Restore Task Report

As you create reports, the documents nest under the node of the relevant report type. When you create a backup job, you can configure a report to display in the body of an email. The report indicates which jobs were run, which succeeded, and which failed. You can also schedule a report to run, save it as a PDF, and send it as an email attachment.

Report information is organized by VM and includes when backups, replicates and restores were run, who ran them, and the status of each. You can create a report for results from one or multiple VMs. Reports can contain customized metrics — for example, date range (day, month), status (successful, failed, stopped). You can run reports as soon as you create them or you can schedule them to be generated later.

My Reports view

The My Reports view is made up of three areas - My Reports on the left, Working Reports on the right, and the Report Template Detail area on the lower right.

My Reports pane

This area features several nodes in a tree structure — Backup Job Summary Report, Backup Task Report, Replicate Job Summary Report, Replicate Task Report, Restore Job Summary Report, and Restore Task Report. Under each of these nodes, all reports that you create and save are nested by type. At the top of the pane, there is a menu bar and toolbar. Within the pane itself, backup and restore reports are designated different icons.

Main toolbar

The icons described in the following table appear in the main toolbar.

Print: Click this icon to print the report displayed in the **Working Reports** pane

Print Preview: use this command to review the report before printing it.

Page Setup: use this command to make adjustments to the report before printing it.

Save: This drop-down menu gives you access to two options — PDF and Excel. When you select a value, the Save As window opens. For the selected report, you can create a PDF or an Excel spreadsheet.

Refresh: This causes recently added/removed reports to be correctly displayed.

Close Report: This closes the reports currently displayed in the Working Reports pane.

Pane toolbar

The icons described in the following table appear in the pane toolbar.

Add: This icon lets you add any of the six reports types using the Restore Wizard:
- Backup Job Summary Report
- Backup Task Report
- Replicate Job Summary Report
- Replicate Task Summary Report
- Restore Job Summary Report
- Restore Task Report

When you first open the **My Reports** view, this is the only icon that is enabled.

Edit: Opens the **Report Wizard**. Configurations for the currently selected report display. You can make changed to the settings and save the s.

Remove: If you click this icon, the report that is currently selected in the **My Reports** pane are removed.

Run: With a reports selected in the **My Reports** pane, click this icon to run it.

Working Reports pane

In this pane, report listings display. The name of the report is at the top of the pane. The format of the text that displays depends on the options you select in the Report Wizard.when you create the report. For example, if you add color or pagination to the report, those settings appear in this pane.

Generating a report using the Report Wizard

Through the Report Wizard, you can generate reports that are filtered by date/time and status. For example, you can generate a report that includes only backup or restore tasks or jobs that failed over the last 10 days.

Complete the steps in the following procedure to schedule a report using the Report Wizard.

To use the Report Wizard

1   In the vRanger interface, click the **My Reports** view.

2   Launch the Report Wizard by doing one of the following:

   - Right-click one of the six nodes, and then click **Add**.

   - Select one of the nodes, and then click the **Add** icon in the view toolbar.

3   On the Report Filter page, enter the information as described in the following table.

Table 10. Available options

Option	Description
Date Range	Narrow report data using these text boxes: • **Over the last:** Enter the number in the first text box and select Days or Months from the drop-down list. Default is 30 days. • **Between:** Click this radio button to enable the two drop-down date lists. Click each to display a calendar from which you can select a date range.
Machine/Source Repository	• **Machine:** Select a machine type — to restrict report data. • **Source:** From this drop-down list, you can select All or any individual source. All licensed, available sources are listed here. • **Repository:** This drop-down list features three values from which to select: All, cifs, or sftp.
Duration	In this section, you can enter a number in the Greater than or equal to __ minutes text box. The value you enter restricts the report to include only jobs that occurred within that time limit. Default is 0.
Job Status	The check boxes available in this section refer to the current status of any job that you might include in reports. Available values are: **Completed**, **Failed**, **Aborted**, and **Cancelled**. NOTE: Because Jobs cannot fail (Tasks within Jobs can), searching for **Job Status - Failed** yields no results. To generate a list of unsuccessful activities, search for **Task Status - Failed**.

4   On the Column Selection page, customize your report by doing the following:

   - Under Selected Columns, select any column you do not want to include in the report, and then click **Clear Column** to move it to the Available Columns text box and remove it from the report. Columns include the following:

       - Job Name

       - Source

       - Repository

       - Job Status

       - Machine Type

       - Start Time

       - End Time

       - Duration (minutes)

       - # Success

       - # Failed

       - # Canceled

       - # Aborted

- To move all columns to the Available Columns text box and select the columns you want to appear in the report, click **Clear All**, select the column you want to include from the Available Columns text box, and then click **Add Column**.

5 On the Header and Footer Style Selection page, select the Body Style, Header Style, and Footer Style as described in the following tables, and then click **Next**.

Table 11. Available Body options

Body style	
Font	Click the Browse (...) button to open the Font dialog; select your preferred Font, Font style, Size, and Effects for the body of the report; and then click **OK**. NOTE: The default font is Arial, 8, Regular.
Even Row Color	To increase readability, select a color in which you want the even rows of the report to appear by clicking the Browse (...), selecting a color, and clicking **OK**.
Odd Row Color	To increase readability, select a color in which you want the odd rows of the report to appear by clicking the Browse (...), selecting a color, and clicking **OK**. NOTE: For best readability, the color selected for the odd rows should be in contrast with the color selected for the even rows.

Table 12. Available Header options

Header style	
Logo	[Optional] Browse to select an image to display in the header.
Text	[Optional] Enter any text that you want displayed in the header.
Font	Click the Browse (...) button to open the Font dialog; select your preferred Font, Font style, Size, and Effects for the report header; and then click **OK**.
Show Time and Date in the report header	Select this option if you want the header to include the current date and time.

Table 13. Available Footer options

Footer style	
Logo	[Optional] Browse to select an image to display in the footer.
Text	[Optional] Enter any text that you want displayed in the footer.
Font	Click the Browse (...) button to open the Font dialog; select your preferred Font, Font style, Size, and Effects for the report footer; and then click **OK**. NOTE: The footer refers to the current status of any job that you might include in reports. Available values are: **Completed**, **Failed**, **Aborted**, and **Cancelled**.
Show Time and Date in the report header	Select this option if you want the footer to include the current date and time.

6 On the Recurrence Schedule Selection page, select a schedule for generating the report or make the report available on demand using the options described in the following table, and then click **Next**.

Table 14. Available options

Option	Description
This will be an On Demand job...	If you select this option, no schedule for the report is set; it occurs on demand. When the page displays, this option is selected by default and all other options are disabled. Clearing this option enables the other options and lets you schedule a report based on a specific time frame.
Start Time	The time of day at which to run the report.

Table 14. Available options

Option	Description

Recurrence Pattern	In this section, you establish how often the report should be generated. There are five options within this section:
	
	• **Daily** – The daily option can be scheduled to run the report every weekday or every *x* number of days.
	
	• **Weekly** – vRanger can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run reports can be configured.
	
	• **Monthly** – The monthly option offers the following configurations:
	
	  • **Day *[x]* of every*[y]* month:**
	
	    • *x* can be any value from 1 to 31. This determines the day of the month on which the job occurs.
	
	    • *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the job to run every other month.
	
	  • **The *[f] [d]* of every *[y]* month(s):**
	
	    • *f* can be either: first, second, third, fourth or last.
	
	    • *d* can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday. *y* can be any value from 1 to 99. This value determines the monthly interval — for example, every two months sets the report to run every other month.
	
	• **Yearly** – The yearly option offers the following configurations:
	
	  • **Every *[m][x]*:**
	
	    • *m* is any month of the year. This determines the month of the year in which the report occurs.
	
	    • *x* can be any value from 1 to 31. This determines the day of the month on which the report occurs.
	
	  • **The *[f][d]* of *[m]*:**
	
	    • *f* can be either: first, second, third, fourth, or last.
	
	    • *d* can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday. *m* is any month of the year. This determines the month of the year in which the report occurs.
	
	• **Interval** – The interval option lets you select the number of days, hours, and minutes that should pass between report jobs.
	
	NOTE: The interval selected must be greater than or equal to five minutes.

Range of Recurrence	Select the start and end date for running the report. • **No end date** – The report continues to run until it is stopped. • **End after** – Select the number of occurrences. • **End by** – Select the date by which to end the recurrence.

7 If you opted to run the report on demand, skip to Step 9.

8 If you scheduled a time for the report to recur, configure the email settings on the Report Publication page described in the following table, and then click **Next**.

Table 15. Available options

Option	Description
To	Use the browse button to select one or more email addresses from the vRanger Address Book to whom you want to send the report. **NOTE:** For more information about adding addresses to the Address Book, see Setting up the SMTP server.
Subject	Enter the text you want to appear in the Subject line of the email that is sent with the report.
Body	Enter the message you want to appear in the body of the email that is sent with the report.
Send as	Select the format in which you want the report to appear. Options include the following: • PDF • Excel

9   On the Report Summary page, verify that the information is correct, and then click **Finish**.

Integrating vRanger

Integrating vRanger

Like any component of your virtual infrastructure, vRanger can be used in conjunction with other applications to achieve more complex tasks.

Performing vRanger sweep-to-tape using Dell™ NetVault™ Backup

Sweeping vRanger backups to tape using NetVault Backup lets you combine vRanger's high-speed backup and recovery capabilities with NetVault Backup's inexpensive, long-term offsite storage technology.

To configure the vRanger sweep-to-tape feature, both vRanger and NetVault Backup must point to a common vRanger repository folder located locally on the vRanger server.

The vRanger backup administrator creates an on-demand (not scheduled), incremental job with a specific name, appending **-Tape** as a case-sensitive suffix (e.g., any-jobname-Tape). The -Tape suffix can be changed by editing the PowerShell script if you prefer something else. There should be a dedicated repository for each on-demand job.

The tape backup software administrator schedules a recurring incremental backup job whose PRE field calls a .cmd file containing the PowerShell script that searches for and then starts the vRanger on-demand job(s) with the -Tape suffix.

The vRanger on-demand job performs the incremental backup, placing the savepoint files in the repository and updating the repository's global manifest file.

After the vRanger job completes (e.g., the job status equals either success or failure), the PowerShell script exits, thus completing the PRE command.

With the PRE command of the tape backup software completed, the tape backup software, which is also pointing to the vRanger repository folder, runs its scheduled incremental job, backing up the vRanger savepoints and global manifest file to tape.

To perform vRanger Sweep-to-Tape using NetVault Backup, complete the following tasks:

- Process of preparing for vRanger sweep-to-tape using NetVault Backup
- Process of performing backup operations
- Performing restore operations

System requirements

For vRanger:

- vRanger on Windows Server 2008, 2008 R2, or 2012

- Virtual or physical machines protected by vRanger

For NetVault Backup:

- A properly configured NetVault Backup Server

- Single Stand-Alone External Tape Drive or a Single Internal Tape Drive in a Library. For example, Dell PowerVault LTO standalone and library-based tape devices are ideal for this solution.

- Windows Server 2008, 2008R2, or 2012

Prerequisite checklist

This guide assumes that both vRanger and NetVault Backup have been installed and configured, and does not cover the installation or the configuration of these applications. Refer to the appropriate documentation links for NetVault Backup and vRanger in the list below for guidance on proper installation and configuration.

- Ranger installed and properly configured.

- Location to store the vRanger repository, accessible by both vRanger and NetVault Backup.

- NetVault Backup installed and properly configured with a tape device or library.

- Blank tape(s) with enough storage to contain the vRanger repository.

Supported devices

Most existing tape libraries and tape formats are supported by NetVault Backup, including stand-alone and library-based tape devices in SAN or shared-SCSI environments. In addition to the line of approved Dell PowerVault tape devices, visit our Supported Devices page for a complete list of supported devices.

System preparation

The following changes need to be made to a default configuration to prepare for the vRanger sweep-to-tape process:

- Enable remote scripts in vRanger PowerShell Console.

- Create a vRanger repository dedicated to sweep-to-tape.

- Download the required vRanger PowerShell scripts.

- Copy the vRanger PowerShell scripts to the NetVault Backup folder.

Process of preparing for vRanger sweep-to-tape using NetVault Backup

The vRanger sweep-to-tape process supports two possible installation scenarios. While both options are similar, the instructions in this guide are designed for the first scenario and can be adapted to fit scenario 2.

Scenario 1: vRanger & NetVault Backup are on the same server:

- Ranger resides on the NetVault Backup Server.

- A local repository shared by vRanger and NetVault Backup.

- PowerShell (PowerShell and vRanger must always be installed on the same server).

- A PowerShell script that finds and launches the vRanger on-demand job.

- A .bat file designed to start the PowerShell script from NetVault Backup Server.

ⓘ | **NOTE:** When installing vRanger and NetVault Backup on the same server do not install the Dell™ NetVault™ Backup Plug-in for VMware®.

Scenario 2: vRanger & NetVault Backup are running on different servers:

NetVault Backup is installed on its own server. All other components reside on the vRanger server.

- **Server 1-NetVault Backup:**
    - NetVault Backup is installed and configured on a supported operating system.
- **Server 2-vRanger/PowerShell:**
    - NetVault Backup agent that can run a PRE command script. The agent runs the PowerShell script contained in a .bat file called by using the PRE command option.
    - vRanger resides on its own server.
    - A local repository located on the vRanger server, and shared by vRanger and NetVault Backup.
    - Windows PowerShell.
    - A PowerShell script that finds and launches the vRanger on-demand job.
    - A command-file script that contains the PowerShell script.

To prepare your environment for vRanger Sweep-to-Tape using NetVault Backup as described in Scenario 1, complete the following tasks:

1 Enabling execution of PowerShell scripts

2 Using vRanger sweep-to-tape scripts

3 Creating a sweep-to-tape vRanger repository

Enabling execution of PowerShell scripts

PowerShell contains an execution policy that determines whether scripts are allowed to run, and whether they must be digitally signed before they can execute. While the default policy is set to block the execution of such scripts, the following section details the process for enabling this option.

Complete the steps in the following procedure to enable execution of PowerShell scripts.

To enable execution of PowerShell scripts

1 Open the vRanger Console, located under **Start > All Programs > Dell Software, Inc. > vRanger Backup & Replication > vRanger Console**.

This opens a PowerShell console ready to accept commands. You should expect to see an error message, which can be remedied by enabling remote scripts.

2 Type the following, and then press **Enter**:

```
Set-executionpolicy remotesigned
```

3 Type **Y** [Enter] to accept the change.

4 Close the vRanger console by typing **Exit** [Enter].

5 Reopen the vRanger console and confirm the proper configuration.

The window scheme changes from red text on a black background to white text on a blue background, and a welcome message appears.

Using vRanger sweep-to-tape scripts

To facilitate starting the vRanger job, Dell has designed scripts that start the vRanger on-demand job from within NetVault Backup.

To use vRanger Sweep-to-Tape scripts

1 Download the vranger_sweep_to_tape.zip.

2 Extract the contents of the zip file to the following folder on your NetVault Backup Server:

C:\Program Files (x86)\Quest Software\NetVault Backup\scripts

NetVault Backup only executes scripts located in this folder.

Creating a sweep-to-tape vRanger repository

For this scenario, you must create a specific vRanger repository to be captured by the NetVault Backup job. This repository should reside in a shared location where both vRanger and NetVault Backup can access the backups. This repository should be used only for vRanger backups that are expected to be swept-to-tape with NetVault Backup.

To create a Sweep-to-Tape vRanger repository

1   Start by opening vRanger from the Start menu:

   Start > All Programs > Dell Software, Inc.> vRanger Backup & Replication > vRanger Backup & Replication

2   From the left-side navigation pane, click **My Repositories**.

3   Right-click your preferred target, select **Add**, and then click the type of storage you want to add.

4   Enter a name for the repository, and then provide the required information and credentials to let vRanger to create and access the repository on this target.

5   Click **OK**.

   vRanger immediately attempts to validate the connection.

Process of performing backup operations

To sweep backed up data to tape, you need to create both a backup job for vRanger and a backup job for NetVault Backup.

To perform backup operations, complete the following tasks:

1   Creating an on-demand vRanger backup job

2   Creating a NetVault Backup job with pre-command scripts

Creating an on-demand vRanger backup job

You need to create a vRanger backup job that captures a complete backup of the machines you want to sweep to tape.

To create an on-demand vRanger backup job

1   Select **My Inventory** in the left-hand navigation pane, and then locate the machine or machines you want to include in the backup.

2   Right-click the computer name and select **Backup [machine name]** on the context menu.

   The Backup Wizard opens.

3   Under Job Name, give the job a name that ends in –Tape; for example, vRanger-Sweep-to-Tape, and then click **Next**.

   ⓘ │ **NOTE:** The provided script is designed to automatically start any vRanger job where the job name
      │ ends with –Tape. Use this naming convention to enable Sweep-to-Tape.

4   On the Virtual Machine Hard Disk Inclusion page, select the disks in the machine to include in the backup job, and then click **Next**.

5   On the repository Selection page, select the target repository created in Creating a sweep-to-tape vRanger repository, and then click **Next**.

6 If the vRanger Transport Selection page appears, select whether to use Automatic transport or to customize the transport method, and then click **Next**.

vRanger Transport determines how backup data is sent and where the restore processing activity occurs. The default transport method is Automatic; however, if the server you are backing up is a VMware virtual machine, you can customize the method as appropriate on the vRanger Transport Selection page.

7 On the Options Selection page of the wizard, select your preferred backup options, and then click **Next**.

8 On the Retention Policy Selection page, determine how many savepoints to keep, and whether to capture Full, Incremental, or Differential backups; and then click **Next**.

> (i) | **NOTE:** Selecting more savepoints provides additional points-in-time to which you can recover the machine. For more information, see About Retention Policies and Space-Saving Technologies.

9 On the Recurrence Schedule Selection page, select **This will be an on demand job**, and then click **Next**.

On-demand jobs are scheduled tasks that are not set to start from a recurring schedule. Because a sweep-to-tape backup job is initiated by script from the NetVault Backup task, you must this job on demand. Doing so saves the job settings but lets the NetVault Backup integration start the job on-demand.

10 On the Email Notifications Selection page, specify who to notify when the job has finished running by entering one or more names and email addresses to be notified, and then click **Next**.

11 On the Summary page, review the Summary window to make sure your configuration is correct, and then do the following:

- Leave the option to **Run the job after Finish is clicked** cleared.
- Click **Finish** to save the job and close the Backup Wizard.

12 To view or access the job, click **My Jobs** in the left-pane, and then select **On Demand jobs** under My Jobs.

Creating a NetVault Backup job with pre-command scripts

Complete the steps in the following procedure to create a NetVault Backup job with pre-command scripts.

To create a NetVault Backup job with pre-command scripts

1 Ensure that NetVault Backup is properly installed, and your tape drive or library is installed and visible to NetVault Backup.

2 Launch the NetVault Backup console, and then click on the Backup icon to start the Backup job configuration process.

3 In the NetVault Backup Backup window, select the preferred NetVault Backup machine from the list of servers, and then double-click it to see a list of identified source devices for this server.

4 Navigate to the location of the vRanger repository, and then select the files and folders to include in the archive.

> (i) | **NOTE:** NetVault Backup supports a wide array of sources from which to capture data, including NAS, CIFS, RAW devices, and local devices.

5 Click the Backup Options tab to customize the backup job.

In this scenario the default settings of Standard Backup Method and Full Backup Type are selected so the job captures a full backup each time the job runs. Additional information on customizing available NetVault Backup options can be found in the *Dell NetVault Backup Administrator's Guide*.

6 Click the Schedule tab to configure a backup interval.

NetVault Backup provides scheduling options such as immediate, one-time, and triggered.

7 Click the Target tab to select the preferred location in which to store the archive, and then do the following:

a Select Specify Device for a list of available devices.

b Select the preferred device in the list to tell NetVault Backup to use this device in the backup.

> ⓘ | **NOTE:** Selecting multiple devices lets NetVault Backup use additional devices as required.

8 Click the Media Option tab to set specific policies to target specific media, label blank media, or specify whether to reuse media with existing data.

9 Under the Advanced Options tab, you can select options to customize the backup process; for example, how long to maintain the archive, encryption, and deduplication options; network compression; and so on. In this example most of the options remain with their default values, with the exception of a Pre Script that executes the vRanger On Demand backup.

10 Under the Pre and Post Scripts section, select **Use Pre Script**, and then enter the following in the entry field:

> vranger_sweep.bat

This process initiates the vRanger backup and waits for the job to complete. After the job is finished, NetVault Backup continues and archives the vRanger repository to tape.

11 In the top right of the window, change the Job Title, as this lets you easily identify this particular backup job for status and reporting purposes; for example, vRanger_Sweep_to_Tape.

12 To submit the job, click the traffic light symbol, which saves and schedules the job with NetVault Backup.

A window opens to indicate the job has been successfully submitted.

13 Close the notification window and the NetVault Backup Backup window.

14 In the NetVault Backup Console, click **Status** to review the current backup status.

The NetVault Backup Server Status Window presents a list of jobs currently running or completed.

15 Under Job Status, right-click the job (for example, vRanger_Sweep_to_Tape), which should have a Run Status of Scheduled, and select **Run Now**, to immediately start the job after asking for confirmation.

This job executes the backup script, which in turn communicates with vRanger and starts the Backup_and_Sweep-Tape job. You can switch back to vRanger to watch the job start and run. After vRanger has finished, NetVault Backup completes the job by capturing the updated repository to tape.

Performing restore operations

There are two available ways to restore your data using vRanger.

To perform restore operations, complete one of the following procedures:

- Restoring a vRanger repository from tape using NetVault Backup

- Restoring from a vRanger repository

Restoring a vRanger repository from tape using NetVault Backup

Use the following steps to export a previously captured vRanger repository from tape.

To restore a vRanger repository from tape using NetVault Backup

1 In the NetVault Backup Console, click **Restore** to open the NetVault Backup Restore Wizard.

2 Under the Selections tab, double-click the server on which the archive is stored.

3 From the list of archives, select the name of the archive (for example, vRanger_Sweep_to_Tape), and then select the folder structure to be restored.

This restores the entire vRanger repository.

> ⚠ **WARNING:** It is important that you restore to a new repository and not to an existing repository. Failure to do so can possibly result in corruption of data.

4 Click the Restore Options tab, and then leave the default settings to maintain the integrity of the vRanger repository.

5 Click the Source tab, and then select the media from which to restore the archive.

6 Click **Target Client**, and select the location to restore the file set from the list of available machines.

7 Click the Schedule tab, and then leave the default schedule option set to **Immediate**.

8 To start the restore process, click the Traffic Light symbol in the menu bar near the top of the window.

9 In the pop-up window, click **OK**.

10 Close the Restore Window and return to the NetVault Backup Console.

11 To review the current backup status, click **Status**, which opens the NetVault Backup Server Status Window and presents a list of jobs currently running or completed.

Restoring from a vRanger repository

The following steps detail the process of restoring data from a vRanger repository.

To restore from a vRanger repository

1 Select the vRanger repository that contains the backup of the virtual machine; for example, Repo-001.

vRanger provides options to restore the complete virtual machine, or just individual files or folders.

A list of protected virtual machines display.

2 Select the virtual machine containing the data to be restored, and then click **File Level Restore** in the top menu to start the process.

A file structure view of the selected virtual machine displays.

3 Locate and highlight the files and folders to be restored to the selected machine.

4 On the right side of the window, select the target location; for example, Desktop.

5 Click **Restore** to start the recovery, which quickly extracts the files from the vRanger repository for immediate access.

After recovery is complete, you now have easy access to the newly restored files.

Deduplication with NetVault SmartDisk and vRanger

NetVault SmartDisk delivers data de-duplication that significantly reduces storage costs. With powerful, byte-level, variable block de-duplication, and advanced data compression that packs more protected data in the same storage area than competing solutions, you can reduce your disk backup storage footprint by up to 90%.

vRanger versions 5.3 and later support the use of a NetVault SmartDisk server as a repository. Simply by adding a NetVault SmartDisk server to vRanger as a repository, backups written to that repository are automatically de-duplicated.

> ⓘ **NOTE:** NetVault SmartDisk repositories are only supported by the Professional edition of vRanger. If you purchased a vRanger Standard license, you cannot use a NetVault SmartDisk repository.

For more information about integrating vRanger and NetVault SmartDisk, see the *Dell vRanger Pro NetVault SmartDisk Integration Guide*.

Data Domain and vRanger repositories

vRanger and Data Domain appliances provide companies with a simple and efficient method for backing up and recovering VMware environments. vRanger provides significant enhancements with regard to architecture, performance and communications over traditional backup solutions. Data Domain provides superior data de-duplication capabilities. The combination of vRanger and Data Domain Technology dramatically reduces your backup and recovery time for virtual machines on VMware ESX hosts.

For more information about integrating vRanger and Data Domain, see the document *vRanger Pro & Data Domain Best Practices*.

Monitoring vRanger

vRanger can be configured to log events to the Windows Event Viewer on the vRanger machine. These events can be filtered and monitored using Microsoft's System Center Operations Manager to monitor replication functions in line with other IT operations. When monitoring is enabled, vRanger creates a vRanger Service node on the local Event Viewer.

The events currently monitored and logged by vRanger include:

- Replication Events – Reserved ID Range 10000 - 10999

  Replication task failure – logged as an error, event ID 10001

- Virtual Appliance Events – Reserved ID Range 11000 - 11999

  VA not powered on – logged as an error, event ID 11001

  VA connection failure – logged as an error, event ID 11002

  VA credential failure – logged as an error, event ID 11003

  VA unknown failure – logged as an error, event ID 11999

  VA return to normal – logged as an Information event, ID 11998

- Datastore Events – Reserved ID Range 12000 - 12999

  Datastore threshold exceeded – logged as an error, event ID 12000

  Datastore unknown failure – logged as an error, event ID 12099

  Datastore return to normal – logged as an Information event, ID 12098

About event monitoring configurations

The configurations for event monitoring are done by way of the configuration file for the vRanger service - *Vizioncore.vRanger.Service.exe.config*. This file is located in the vRanger installation directory. By default this is:

```
C:\Program Files\Dell\vRanger
```

⚠ | **WARNING: Change only the values listed below unless directed otherwise by Dell Support.**

To configure event monitoring:

Opening the configuration file

Complete the steps in the following procedure to open the configuration file.

To open the configuration file

1   Open the *Vizioncore.vRanger.Service.exe.config* with your text editor of choice.

2   Carefully make the required changes, using the information below as guidance.

> ⓘ **NOTE:** Not all text editors include line numbers. If you do not see the line numbers referenced below in your text editor, this is not an indication of a problem.

Configuring replication monitoring

Replication monitoring logs replication task failures to the Windows event log with the Event ID of 10001. The log also includes the error message and information about the task such as source virtual machine, target host, and virtual center.

To configure replication monitoring

1   In the *Vizioncore.vRanger.Service.exe.config* file, scroll to line 162 or search for **"Replication FailureWriteEventLog"**.

2   Change the value "false" to "true", ensuring the quotation marks remain.

3   Save the file and exit.

Continue to the next section to make additional changes.

Configuring virtual appliance monitoring

When enabled, at the specified interval, the vRanger Virtual Appliance (VA) monitor goes through all of the configured VAs and checks to make sure that their power is turned on and that vRanger is able to connect to them using the entered IP addresses and credentials.

To configure virtual appliance monitoring

1   In the *Vizioncore.vRanger.Service.exe.config file*, find the lines corresponding to the preferred setting, as shown below.

- **VirtualApplianceMonitoringEnabled** – line 164.

  Set this value to "true" to enable monitoring of the virtual appliance.

- **VirtualApplianceMonitoringEventLogFailureLogMax** – line 167.

  This is the maximum number of times that a particular VA failure can be logged to the Event log before a success. Change the value of "7" to the preferred value, ensuring the quotation marks remain.

  This setting is intended to limit the number of messages in your event viewer. A refresher message appears periodically based on the value below.

- **VirtualApplianceMonitorSecondaryIntervalInHours**

  This value determines how often the refresher message occurs for a failed virtual appliance check after the FailureLogMax value has been reached. The default value is "6", which indicates that it logs an event every six hours.

- **VirtualAppliancesConnectionCheckIntervalInMinutes**

  This setting determines how often vRanger checks the connection to configured virtual appliances. The default value is "5", which indicates that it checks the connection every five minutes.

2   Save the file and exit.

Continue to the next section to make additional changes.

Configuring datastore monitoring

When enabled, at the specified interval, the datastore monitor goes through all of the configured datastores and determines if the space available on the datastore exceeds the specified threshold. This may be either by percentage or a specific value in gigabytes. When this occurs, an event is logged in the windows event log with an ID of 12000.

To configure datastore monitoring

1 In the *Vizioncore.vRanger.Service.exe.config file*, find the lines corresponding to the preferred setting, as shown below.

   • **DatastoreMonitoringEnabled** – line 171

      Set this value to "true" to enable monitoring of the datastores.

   • **DatastoreMonitoringSpaceThresholdValue** – line 173

      This is the threshold value for a datastore space before a warning is sent. This value is modified my the **DatastoreMonitoringSpaceThresholdType** setting below.

   • **DatastoreMonitoringSpaceThresholdType** – line 175

      This value determines whether vRanger checks for percentage of free space or an absolute value un GB. Set "0" for percent, "1" for GB.

   • **DatastoreMonitorCheckIntervalInMinutes** – line 177

      How often vRanger checks the datastores for low space (in minutes). Default value: 30 minutes.

   • **DatastoreMonitoringEventLogFailureLogMax** – line 180

      The maximum number of times a particular datastore monitor failure can be logged to the windows event log before a success. The default value is "7".

      This setting is intended to limit the amount of messages in your event viewer. A refresher message appears periodically based on the value below.

   • **DatastoreMonitorSecondaryIntervalInHours** -

      This value determines how often the refresher message occurs for a datastore that has reached its space threshold value after the FailureLogMax value has been reached. The default value is "6", which indicates it logs an event every six hours.

2 Save the file and exit.

Restart the vRanger service

In order for the changes to take effect, you must restart the vRanger service.

To restart the vRanger service

1 On the vRanger machine, click **Start**, then Run. Type "services.msc" and click **OK**.

2 Find the vRanger Service. Right-click on the service, and click **restart**.

Using the vRanger Console

- vRanger Console overview
- Configuring the vRanger Console
- vRanger vAPI Cmdlets

vRanger Console overview

vRanger uses Microsoft PowerShell with a Dell snap-in to provide scripting access to vRanger features and functions. This lets vRanger tightly integrate with other data protection solutions and offers an efficient configuration option for advanced users.

The snap-in includes the vAPI cmdlets required to create and manage jobs as well as view relevant information data reading your virtual environment.

For more information on Microsoft Powershell, refer to Windows Powershell Owner's Manual:
http://technet.microsoft.com/en-us/library/ee177003.aspx

Configuring the vRanger Console

The vRanger installer automatically installs Microsoft PowerShell 1.0 and the vAPI snap-in during the product installation. Before using the vAPI cmdlets, you must first configure the PowerShell execution policy to run the vRanger scripts. This is a one-time configuration that does not need to be repeated.

To configure the vRanger Console

1   From the **Start** menu, navigate to the vRanger PowerShell directory.

- By default this is: **Start > All Programs > Dell Software, Inc. > vRanger Console**

2   Select **vRanger Console**.

When opening the vRanger Console for the first time, you see an error message stating that the execution of scripts is disabled on your system. This is due to the default PowerShell execution policy, which must be enabled.

3   At the prompt, type the command below and press **Enter**.

```
Set-ExecutionPolicy RemoteSigned
```

> (i) | **NOTE:** PowerShell and vAPI commands are not case sensitive.

4   When the prompt re-appears, close and re-open the vRanger Console.

When the vRanger Console restarts, you should see the screen below, with a blue background and a list of vAPI cmdlets.

The vRanger vAPI is now configured and ready for use.

vRanger vAPI Cmdlets

The vRanger PowerShell snap-in includes the cmdlets listed below.

Inventory management

Add-DdbRepository

Add-PhysicalMachine

Add-VirtualAppforLinuxFLR

Add-VirtualAppforLinuxFLRVA

Add-VirtualCenter

Get-Datastore

Get-InventoryEntities

Get-IsInventoryRefreshing

Get-VirtualApplianceConfig

Get-VirtualApplianceDeploymentStatus

Get-VirtualApplianceReconfigStatus

Get-VirtualMachinesUnderInventory

Remove-VirtualAppliance

Remove-VirtualApplianceConfiguration

Update-HypervRestoreJobTemplate

Update-VirtualAppliance

Update-VirtualApplianceConfiguration

Configuration

Dismount-LinuxVolume

Get-MonitorLog

Get-VirtualApplianceConfig

Get-VmDisk

Get-VMDKVolume

Install-VirtualAppliance

Mount-LinuxVolume

New-EmailAddress

New-SMTPServer

New-TransportConfiguration

New-VirtualAppliance

Remove-AllMount

Remove-LinuxVolume

Set-CBTonVM

Set-LinuxVolume

Set-MountPath

Set-Resources

Job management

Disable-Job

Enable-Job

Get-AddressBook

Get-ConfigOption

Get-Connection

Get-CurrentTemplateVersionID

Get-Job

Get-JobTemplate

New-Daily Schedule

New-IntervalSchedule

New-MonthlySchedule

New-WeeklySchedule

New-YearlySchedule

Remove-JobTemplate

Remove-VirtualApplianceConfiguration

Set-Cataloging

Stop-vRangerJob

Repositories

Add-CIFSRepository

Add-DdbRepository

Add-FTPRepository

Add-HypervCluster

Add-NVSDRepository

Add-SFTPRepository

Get-PhysicalMachineDiskMap

Get-RepositoryJob

Get-RepositorySavePoint

Get-Savepoints

Remove-Repository

Remove-SavePoint

Remove-DdbStorageUnit

Backup

Add-BackupGroupEntity

Add-BackupJobTemplate

Get-BackupGroupEntity

New-BackupFlag

New-BackupGroupMember

Remove-BackupGroupEntity

Remove-BackupGroupMember

Update-BackupJobTemplate

Restore

Add-RestoreFromManifestJobTemplate

Add-RestoreJobTemplate

Get-CatalogSearchData

Get-CatalogStatus

Get-RestoreDiskMap

Get-SavepointManifest

New-RestoreFlag

Remove-Catalog

Update-RestoreJobTemplate

Replication

Add-PhysicalMachine

Get-GlobalTransportFailover

Get-MonitorLog

Get-TransportFailover

New-ReplicationFlag

Run-ReplicationFailover

Run-ResumeReplicationFailover

Run-TestReplicationFailover

Update-GlobalTransportFailover

Update-ReplicationJobTemplate

Cmdlet help

To get help on any of the vAPI cmdlets, use the Get-Help command. At the prompt, enter Get-Help *command*, where *command* is any of the above vAPI cmdlets.

For example, enter **Get-Help Add-SFTP Repository** and the following is displayed:

```
Description The Add-SftpRepository cmdlet adds a SFTP repository to vRanger. The
repository can be later used to store backup files from backup jobs.

Syntax Add-SftpRepository [-Server] <string> [[-Port] [<int>]]
[[-TargetDirectory][<string>]] [-User] <string> [-Password] <string> [[-TimeOut]
[<short>]] [-Name] <string> [[-Description] [<string>]] [[-Overwrite]
[<Boolean>]][<CommonParameters>]
```

> ⓘ **NOTE:** PowerShell and vAPI commands are not case sensitive. vAPI cmdlets are shown capitalized only to make them easier to read.

For more information, you may use the "-detailed" or "-full" parameters. For example:

Get-Help Add-SFTPRepository -Detailed or **Get-Help Add-SFTP Repository - Full**.

vAPI Cmdlet details

> ⓘ **IMPORTANT:** For the most up-to-date help contents, refer to the commandlet help. To get help on any of the vAPI cmdlets, use the Get-Help command. At the prompt, enter:
>
> `Get-Help command -full`
>
> where `command` is any of the above vAPI cmdlets.

Add-BackupGroupEntity

Description	Adds a new Backup Group to vRanger. New-BackupGroupMember cmdlet is used to add Members to a Backup Group created using this cmdlet or the UI.
Syntax	`Add-BackupGroupEntity [-BackupGroupName] <String> [<CommonParameters>]`
Parameters	**BackupGroupName** <String> – The BackupGroup Name **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type "get-help about_commonparameters". **Type [-Type]** – This optional parameter specifies the type of backupgroup: VMware, HyperV or Physical. If not specified, the backupgroup defaults to VMware.
Examples	`Add-BackupGroupEntity -BackupGroupName BG1`
Related Links	Remove-BackupGroupEntity New-BackupGroupMember Remove-BackupGroupMember Get-BackupGroupEntity

Add-BackupJobTemplate

Description	Adds a backup job. The user also has the option of running the job immediately after the template is created.
Syntax	`Add-BackupJobTemplate [-JobName] <String>`

```
Add-BackupJobTemplate [-JobName] <String>
[[-JobDescription] <String>] [-JobEntity] <Entity>
[[-ExcludedVMList] <String>] [[-IncludedDiskList] <String>] [-
TargetRepository] <Repository>
[[-JobSchedule] <Schedule>] [-Flags] <NoFlags | BackupPoweredOnOnly
| CheckDestinationFreeSpace | UseCompression | UpdateNotes |
EnableGuestQuiescing | EnableABM | UseFiber |
PerformNetworkBackupOnFailure | UseCataloging | AllFlags> [-
NumberOfSavePoints] <Int> [-SpaceSavingTechnologyTypeFlag] <None |
Incremental | Differential>
[[-SpaceSavingCountThreshold] <Int>]
[[-SpaceSavingPercentSizeThreshold] <Int>]
[[-NotificationList] <String>] [[-RunJobNow] <Boolean>] [[-
TransportConfiguration]
<TransportConfiguration>] [[-FullBackupDays]
<List`1>][<CommonParameters>]
```

Parameters	- **JobName** <String> – Job name
	- **JobDescription** [<String>] – Job description
	- **JobEntity** <Entity> – Inventory entity on which the job is going to run on — ESX host, VirtualCenter etc.
	- **ExcludedVMList** [<String>] – List of virtual machine UUIDs to exclude from this job
	- **IncludedDiskList** [<String>] – List of virtual disks to include in this job
	- **TargetRepository** <Repository> – Repository where the backups are to be stored
	- **JobSchedule** [<Schedule>] – Schedule on which the job runs
	- **Flags** <BackupFlags> – Backup flags
	- **NumberOfSavePoints** <Int> – Max number of savepoints to st [ore before older ones are removed.
	- **SpaceSavingTechnologyTypeFlag** <SpaceSavingTechnologyType> – Spacesaving technology to use: differential, incremental or none.
	- **SpaceSavingCountThreshold** [<Int>] – The number to save before performing a full backup.
	- **SpaceSavingPercentSizeThreshold** [<Int>] – The size to allow as a % of the virtual machine hard disk file before performing a differential backup.
	- **NotificationList** [<String>] – Email notification list for the job.
	- **RunJobNow** [<Boolean>] – Whether to run the job now
	- **TransportConfiguration** [<TransportConfiguration>] – The transport configuration options to use for this job.
	- **FullBackupDays** [<List '1'>] – The days to perform a full backup.
	- **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
| Examples | ```
Add-BackupJobTemplate -JobName myjob -JobEntity $inv -
TargetRepository $repo -JobSchedule $s -Flags NoFlags -
NumberOfSavePoints 3 -SpaceSavingTechnologyTypeFlag None
``` |
	This command adds a backup job template. Because the RunJobNow flag is not specified, the job does not run immediately.
Related Links	Add-RestoreJobTemplate
	Add-PhysicalMachine
	Update-BackupJobTemplate
	Get-Job
	Get-JobTemplate

Add-CIFSRepository

Description	The Add-CIFSRepository cmdlet adds a CIFS repository to vRanger. The repository can be later used to store backup files from backup jobs.
| Syntax | ```
Add-CifsRepository [-Server] <String> [-ShareName] <String> [[-
TargetDirectory] <String>] [-DomainName] <String> [-User] <String>
[-Password] <String> [-Name] <String> [[-Description] <String>] [[-
Overwrite] <Boolean>] [[-EnableEncryption] <Boolean>]
[[-EncryptionPassword] <String>] [[-ImportOption]

<String>] [<CommonParameters>]
``` |

Parameters	• **Server** \<String\> – The CIFS Server FQDN or IP. • **ShareName** \<String\> – Name of the file system share the repository is on • **TargetDirectory** [\<String\>] – The CIFS Target Directory. • **DomainName** \<String\> – Name of the domain • **User** \<String\> – The CIFS User Name. • **Password** \<String\> – The CIFS Password. • **Name** \<String\> – A friendly name describing this CIFS server. • **Description** [\<String\>] – The Description describing the CIFS server. • **Overwrite** [\<Boolean\>] – Whether to overwrite the existing repository if one exists • **EnableEncryption** [\<Boolean\>] – Whether to enable encryption • **EncryptionPassword** [\<String\>] – The encryption Password. • **ImportOption** [\<String\>] – Whether to import the save points as read-write or read-only. • **\<CommonParameters\>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-CifsRepository -Server myserver.dev.corp -Sharename myshare -TargetDirectory myrepo -User myname -Password mypassword -Name myCifs Repo -Overwrite 1` This command adds a repository on server myserver.dev.corp pointing to a share named myshare using the credentials provided. Existing repositories at the location specified are overwritten.
Related Links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-HypervCluster Add-NVSDRepository Add-SFTPRepository Add-FTPRepository

Add-DdbReplicationRepository

Description	
Syntax	`Add-DdbReplicationRepository [-Server] <String> [StorageUnit] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>][-RepositoryId] <String> [-Verbose] [-Debug] [-ErrorAction <ActionPreference>] [-WarningAction <ActionPreference>] [-ErrorVariable <String>] [-WarningVariable <String>] [-OutVariable <String>][-OutBuffer <Int32>] [-WhatIf] [-Confirm]`
Parameters	
Examples	
Related Links	

Add-DdbRepository

Description	Adds a DDB repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	`Add-DdbRepository [-Server] <String> [-StorageUnit] <String> [-User] <String>[-Password] <String> [-Name] <String> [[-Description] <String>][[-Overwrite] <Boolean>] [[-ImportOption] <String>] [<CommonParameters>]`
Parameters	• **Server** <String> – The DDB Server FQDN or IP. • **StorageUnit** <String> – The DDB Storage Unit • **User** <String> – The User Name to access the DDB Server. • **Password** <String> – The Password to access the DDB Server. • **Name** <String> – A friendly name describing this DDB server. • **Description** [<String>] – The Description describing the DDB server. • **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists • **ImportOption** [<String>] – Whether to import the save points as read-write or read only. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-DdbRepository -Server myserver.dev.corp -StorageUnit mystorage -User myname -Password mypassword -Name myDDB Repo -Overwrite 1` This command adds a repository on server myserver.dev.corp pointing to storage unit named mystorage using the credentials provided. Existing repositories at the location specified are overwritten.
Related Links	Get-PhysicalMachineDiskMap Add-CIFSRepository Add-HypervCluster Add-NVSDRepository Add-SFTPRepository Add-FTPRepository

Add-EsxHost

Description	Adds an ESX host to the inventory
Syntax	`Add-EsxHost [-DNSorIP] <String> [-Username] <String> [-RootPassword] <String> [[-UserPassword] <String>] [[-Port] <Int>] [[-AddUser] <Boolean>] [<CommonParameters>]`
Parameters	• **DNSorIP** <String> – The DNS or IP for the ESX/ESXi Host. • **Username** <String> – The Username for the ESX/ESXi Host. • **RootPassword** <String> – The Root password for the ESX/ESXi Host. • **UserPassword** [<String>] – The User password for the ESX/ESXi Host. • **Port** [<int>] – The SSH Port for the ESX/ESXi Host. • **AddUser** [<Boolean>] – Add user of the user does not exist. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Add-EsxHost -DNSorIP 10.6.20.20 -Username root -RootPassword`
	`p@ssw0rd`
	
	This command adds an EsxHost to the inventory.

Related Links	Add-VirtualCenter Get-InventoryEntities

Add-FTPRepository

Description	The Add-FtpRepository cmdlet adds a FTP repository to vRanger. The repository can be later used to store backup files from backup jobs.

Syntax	`Add-FtpRepository [-Server] <String> [[-Port] <Int32>] [-`
	`TargetDirectory] <String> [-User] <String>`
	`[-Password] <String> [[-TimeOut] <Int16>] [-Name]`
	`<String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-`
	`EnableEncryption] <Boolean>]`
	`[[-EncryptionPassword] <String>] [[-ImportOption] <String>]`
	`[<CommonParameters>]`

Parameters	• **Server** <String> – The Ftp Server FQDN or IP. • **Port** [<int>] – The FTP Server Port. • **TargetDirectory** [<String>] – The FTP Target Directory. • **User** <String> – The FTP user name. • **Password** <String> – The FTP Password. • **TimeOut** [<short>] – Timeout value. • **Name** <String> – A friendly name describing this Ftp server. • **Description** [<String>] – The Description describing the Ftp server. • **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists. • **EnableEncryption** [<Boolean>] – Whether to enable Encryption. • **EncryptionPassword** <String> – The Encryption Password. • **ImportOption** [<String>] – Whether to import the save points as read-write or read-only. Possible options 'ReadOnly' and 'ReadWrite'. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Add-FtpRepository -Server myserver.dev.corp -TargetDirectory mydir`
	`-User myname -Password mypassword -Name myrepo -Overwrite 1`
	
	This command adds a repository on server myserver.dev.corp pointing to storage unit named mystorage using the credentials provided. Existing repositories at the location specified are overwritten.

Related Links	Get-PhysicalMachineDiskMap Add-CIFSRepository Add-HypervCluster Add-NVSDRepository Add-SFTPRepository

Add-HypervCluster

Description	Adds a Hyper-V Cluster to the inventory.
Syntax	`Add-HypervCluster [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int32>] [[-IsInstallAgent] <Boolean>] [<CommonParameters>]`
Parameters	• **DNSorIP** <String> – The DNS or IP for the Hyper-V Cluster • **Username** <String> – The Username to add the Hyper-V Cluster • **Password** <String> – The Password to add the Hyper-V Cluster • **Port [**<Int32>] – The Port to add the Hyper-V Cluster • **IsInstallAgent** [<Boolean>] – Install agent on Cluster • **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-HypervCluster -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd` This command adds a Hyper-V Cluster to the inventory.
Related Links	Add-BackupJobTemplate Get-InventoryEntities

Add-HypervHost

Description	Adds a Hyper-V Host to the inventory.
Syntax	`Add-HypervHost [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int32>] [[-IsInstallAgent] <Boolean>] [<CommonParameters>]`
Parameters	• **DNSorIP** <String> – The DNS or IP for the Hyper-V Host • **Username** <String> – The Username to add the Hyper-V Host • **Password** <String> – The Password to add the Hyper-V Host • **Port** [<Int32>] – The Port to add the Hyper-V Host • **IsInstallAgent** [<Boolean>] – Install agent on Host • **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-HypervHost -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd` This command adds a Hyper-V Host to the inventory.
Related Links	Add-BackupJobTemplate Get-InventoryEntities

Add-HypervRestoreJobTemplate

Description	Adds a Hyper-V restore job template. Users also have the option of running the job immediately after the template is created.
Syntax	`Add-HypervRestoreJobTemplate [-JobName] <String> [[-JobDescription] <String>] [-SavePoint] <SavePoint> [[-JobSchedule] <Schedule>] [[-Flags] <NoFlags| ForceOverwrite	ForcePowerOn	UseFiber	PerformNetworkRestoreOnFailure	AddHypervVirtualMachineToCluster	AllFlags>] [[-NotificationList] <List `1] [[-RunJobNow] <Boolean>] [[-VMName] <String>] [[-TargetHost] <String>][[-VmLocation] <String>] [[-SkipDisks] <List`1>] [[-DiskLocations] <List`1>] [[-TargetNetworks] <List`1>] [[-SkipAdapters] <List`1>] [<CommonParameters>]`
Parameters	**JobName** <String> – Job name **JobDescription** [<String>] – Job description **SavePoint** <SavePoint> – Savepoint to restore **JobSchedule** [<Schedule>] – Schedule on which the job runs. If not set, the job runs on demand. **Flags** [<RestoreFlags>] – Restore flags. If not set, the job has no flags. **NotificationList** [<List`1>] – Email addresses separated by comma. **RunJobNow** [<Boolean>] – Whether to run the job now. If not set, the job does not run immediately after creation of the template. **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used **TargetHost** [<String>] – Fully qualified domain name or IP address of the target host. **VmLocation** [<String>] – Location on target host where the restored VM is located. If not set, the default value for the host is used. **SkipDisks** [<List`1>] – Sequence number (base 0) of the disks to skip from the savepoint. Multiple numbers are separated by comma.</li><li>**DiskLocations** [<List`1>] – Disk locations to use on the target host. Multiple values are separated by comma. If only one is supplied, it is used for all disks on the VM. If not set, the default disk location for the host is used for all disks. **TargetNetworks** [<List`1>] – Networks to use on the target host. Multiple values are separated by comma. If only one network name is provided, it is used for all adapters on the VM. If not set, the original networks are used when restoring to the original VM.</li><li>**SkipAdapters** [<List`1>] – Sequence number (base 0) of the adapters to skip in the restore job. Multiple numbers are separated by comma. **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".

Examples

------------------EXAMPLE 1-------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$spList = get-repositorysavepoint $repo.id
$mysp = ($spList)[1]
Add-RestoreJobTemplate -Jobname vapirestore  -Savepoint $mysp
```

This command creates a restore job template that restores to the original VM in the savepoint. The job runs on demand.

------------------EXAMPLE 2-------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type HypervHost |
where{$_.name -eq 'dev 015.dev.eng.int'}
$targetDiskLocation = 'c:\hyper-virtual hard disks
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
Add-RestoreJobTemplate -jobname vapirestore_totaraget  -savepoint
$mysp -targethost $targethost.name -DiskLocations
$targetDiskLocation -VmLocation 'c:\hyper-v\configuration' -
targetnetworks $targetnw.name
```

This command creates a restore job template that restores to the specified target host. It also specifies disk location, vm location, and the network to use on the target host.

------------------EXAMPLE 3-------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type HypervHost |
where{$_.name -eq 'dev015.dev.eng.int'}
$targetDiskLocation = 'c:\hyper-virtual hard disks
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$skipdisks = '0'
Add-RestoreJobTemplate -jobname vapirestore_skipdisk  -savepoint
$mysp -targethost $targethost.name -DiskLocations
$targetDiskLocation -VmLocation 'c:\hyper-v\configuration' -
targetnetworks $targetnw.name -SkipDisks $skipdisk s
```

This command creates a restore job template that restorse to the specified target host. The first disk (0) in the savepoint is skipped by the restore job.

Related Links

Add-BackupJobTemplate

Add-RdaRepository

Update-RestoreJobTemplate

Get-Job

Get-JobTemplate

Add-NFSRepository

Description	The Add-NfsRepository cmdlet adds a Network File Share repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	`Add-NfsRepository [-Server] <String> [-ExportDirectory] <String> [-TargetDirectory] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite][<Boolean>]] [[-EnableEncryption] [<Boolean>]] [[-EncryptionPassword] <String>] [[-BlockSize] <Int>] [[-EnableCompress] [<Boolean>]] [[-ImportOption] [<string>]] [<CommonParameters>]`
Parameters	**Server** <String> – The NFS Server FQDN or IP. **ExportDirectory** <String> – The NFS Export Directory. **TargetDirectory** <String> – The NFS Target Directory. **Name** <String> – A friendly name describing this NFS server. **Description** <String> – The Description describing the NFS server. **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists **EnableEncryption** [<Boolean>] – Whether to enable encryption on backups on this repository **EncryptionPassword <String>** – The encryption Password. **ImportOption** [<String>] – Whether to import the save points as read-write or read only. Possible options 'ReadOnly' and 'ReadWrite'. **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Add-NfsRepository -Server myserver.dev.corp -ExportDirectory /vol/vol1 -TargetDirectory /test -Name MyNFSRepo -Overwrite 1`
	
	This command adds a repository on server myserver.dev.corp. Existing repositories at the location specified are overwritten.
	
	-------------------**EXAMPLE 2**-------------------
	
	`Add-NfsRepository -Server myserver.dev.corp -ExportDirectory /vol/vol1 -TargetDirectory /test -Name MyNFSRepo -EnableEncryption 1 -EncryptionPassword myPassword -EnableCompress 1`
	
	This command adds a repository on server myserver.dev.corp. Enables encryption on backups of this repository. All backups to this repository are compressed.

Related Links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-NVSDRepository Add-CIFSRepository Add-SFTPRepository Add-FTPRepository

Add-NVSDRepository

Description	The Add-NvsdRepository cmdlet adds a NetVault SmartDisk repository to vRanger. The repository can be later used to store backup files from backup jobs.

Syntax	`Add-NvsdRepository [-Server] <String> [-TargetDirectory] <String> [-User] <String>`
	`[-Password] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-ImportOption] <String>]`
	`[<CommonParameters>]`

Parameters	• **Server** <String> – The NetVault SmartDisk server FQDN or IP. • **TargetDirectory** [<String>] – The NetVault SmartDisk Target Directory. • **User** <String> – The NetVault SmartDisk user name. • **Password** <String> – The NetVault SmartDisk Password. • **Name** <String> – A friendly name describing this NetVault SmartDisk server. • **Description** <String> – The Description describing the NetVault SmartDisk server. • **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists • **ImportOption** [<String>] – Whether to import the save points as read-write or read only. Possible options 'ReadOnly' and 'ReadWrite'. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-NvsdRepository -Serve r myserver -User myname -Password mypassword -Name reponame -tagetDirectory /targetDirectoryName -Overwrite 1` This command adds a repository on server myserver pointing to a target directory named "targetDirectoryName" using the credentials provided. Existing repositories at the location specified are overwritten.
Related Links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-HypervCluster Add-CIFSRepository Add-SFTPRepository Add-FTPRepository

Add-PhysicalMachine

Description	Adds a physical machine to the inventory.
Syntax	`Add-PhysicalMachine [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int>]` `[-IsInstallAgent] <Boolean>] [[-AgentLocation] <String>] [<CommonParameters>]`

Parameters	• **DNSorIP** <String> – The DNS or IP for the Physical Machine • **Username** <String> – The Username to add the Physical Machine • **Password** <String> – The Password to add the Physical Machine • **Port** [<int>] – The Port to add the Physical Machine • **IsInstallAgent** [<Boolean>] – Install Agent on Machine • **AgentLocation** [<String>] – Agent Installation Location • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug,-ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-PhysicalMachine -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd` This command adds a replication job template.
Related Links	Add-BackupGroupEntity Get-InventoryEntities

Add-RdaRepository

Description	The Add-RdaRepository cmdlet adds an RDA repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	`Add-RdaRepository [-Server] <String> [-StorageUnit] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-ImportOption] <String>] [<CommonParameters>]`
Parameters	• **Server** <String> – The DR server FQDN or IP. • **StorageUnit** [<String>] – The RDA logical storage unit. • **User** <String> – The user name to access the DR server. • **Password** <String> – The Password to access the DR server. • **Name** <String> – A friendly name describing this DR server. • **Description** <String> – The Description describing the NetVault SmartDisk server. • **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists • **ImportOption** [<String>] – Whether to import the save points as read-write or read only. Possible options 'ReadOnly' and 'ReadWrite'. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-RdaRepository -Server myserver -StorageUnit mystorage -User myname -Password mypassword -Name reponame -targetDirectory /targetDirectoryName -Overwrite 1` This command adds a repository on server myserver pointing to a target directory named "targetDirectoryName" using the credentials provided. Existing repositories at the location specified are overwritten.
Related Links	Get-Repository Add-DdbRepository Add-CIFSRepository Add-SFTPRepository Add-FTPRepository Add-NVSDRepository

Add-ReplicationJobTemplate

Description	Adds a replication job template. The user also has the option of running the job immediately after the template is created.
Syntax	`Add-ReplicationJobTemplate [-JobName] <String>` `[[-JobDescription] <String>] [-JobEntity] <VMEntity>` `[-TargetHost] <HostEntity> [-ReplicateName] <String> [-Type]` `<String> [[-TargetNetworks] <String>]` `[[-PrimaryDatastore] <String>] [[-TargetDatastores] <String>] [[-` `JobSchedule] <Schedule>]` `[[-TransportFailover] <String>]` `[[-NotificationList] <String>] [-Flags]<NoFlags |` `ReplicatePoweredOnOnly	CheckDestinationFreeSpace	UseCompression` `| UpdateNotes	EnableGuestQuiescing	EnableABM	AllFlags> [[-` `RunJobNow] <Boolean>]` `[[-IncludedDisk] <String>] [<CommonParameters>]`
Parameters	• **JobName** <String> – Job name. • **JobDescription** [<String>] – Job description. • **JobEntity** <VMEntity> – VM entity on which the job runs. • **TargetHost** <HostEntity> – Target Host entity to which the VM is replicated • **ReplicateName** [<String>] – Replicated VM Name • **Type** <String> – Replication type. Valid types are Differential and Hybrid • **TargetNetworks** [<String>] – Networks to use on the target host. Multiple values are separated by a comma. • **PrimaryDatastore** <String> – Primary Datastore on target host where the replicated VM are to be located. • **TargetDatastores** [<String>] – Datastores to use on the target host. Multiple values are separated by a comma. • **JobSchedule** <Schedule> – The recurrence schedule on which the job runs. • **TransportFailover** [<String>] – Transport Failover order separated by ',' to be attempted until the task is successful. Valid types are VA-HotAdd, VA-Network, COS. • **NotificationList** [<String>] – Email notification list for the job. Separate multiple emails by ','. • **Flags** <ReplicationFlags> – Replication flags: NoFlags, ReplicatePoweredOnOnly, CheckDestinationFreeSpace, UseCompression, UpdateNotes, EnableGuestQuiescing, AllFlags. Multiple flags can be separated by '|' • **RunJobNow** [<Boolean>] – Whether to run the job immediately after the job template is created. • **IncludedDisk** [<String>] – List of disks to include in this job. All VM disks are replicated if none is specified. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	------------------EXAMPLE 1------------------

```
Add-ReplicationJobTemplate -JobName vRepJob1 -JobEntity $vm -
TargetHost $targethost -ReplicateName TestRep -Type Differential -
Flags $flags -targetDatastores MyDatastore1, MyDatastore2
-PrimaryDatastore MyDatastore -targetnetworks VMNetwork -
TransportFailover VA-HotAdd, VA-Network -RunJobNow 1 -includeddisk
0,2,4
```

This command adds a replication job template.

------------------EXAMPLE 2------------------

```
Add-ReplicationJobTemplate -JobName vRepJob1 -JobEntity $vm -
TargetHost $targethost -ReplicateName TestRep -Type Differential -
Flags $flags -targetDatastores MyDatastore1, MyDatastore2
-PrimaryDatastore MyDatastore -targetnetworks VMNetwork -
includeddisk 0,1,2
```

This command adds a replication job template. Because the RunJobNow flag is not specified, the job does not run immediately. Because no transport failover is specified, the default order of transport failovers are used.

Related Links	Add-BackupJobTemplate
	Add-RestoreJobTemplate
	Update-ReplicationJobTemplate
	Get-Job
	Get-JobTemplate

Add-RestoreFromManifestJobTemplate

Description	Adds a restore job template from a manifest file. The manifest file can be a vRanger manifest file or a Legacy vRanger manifest file. Users also have the option of running the job immediately after the template is created.
Syntax	`Add-RestorefromManifestJobTemplate [-JobName] <String>` `[[-JobDescription] < String>] [-Username] <String> [-Password]` `<String> [-ManifestPath] <String>` `[[-EncryptionPassword] <String>] [[-JobSchedule] <Schedule>] [[-` `Flags]<No Flags	ForceOverwrite	ForcePowerOn	UseFiber |` `PerformNetworkRestoreOnFailure	AllFlags>]` `[[-NotificationList] <String>] [[-RunJobNow] <Boolean>][[-VMName]` `<String>] [[-TargetHost] <String>][[-PrimaryDatastore] <String >]` `[[-SkipDisks] <String>] [[-TargetDatastores] <String>]` `[[-TargetNetworks] <String>] [[-SkipAdapters] <String>][[-` `TransportConfiguration] <TransportConfiguration>]` `[<CommonParameters>]`

Parameters	• **JobName** <String> – Job name

Parameters

- **JobName** <String> – Job name
- **JobDescription** [<String>] – Job description
- **Username** <String> – Username
- **Password** <String> – Password
- **ManifestPath** <String> – Path for the manifest file
- **EncryptionPassword** [<String>] – Encryption Password
- **JobSchedule** [<Schedule>] – Schedule on which the job runs. If not set, the job runs on demand.
- **Flags** [<RestoreFlags>] – Restore flags. If not set, the job has no flags.
- **NotificationList** [<String>] – Email addresses separated by comma.
- **RunJobNow** [<Boolean>] – Whether to run the job now. If not set, the job does not run immediately after creation of the template.
- **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used
- **TargetHost** [<String>] – Fully qualified domain name or IP address of the target host.
- **PrimaryDatastore** [<String>] – Datastore on target host where the restored VM is located. If not set, the original primary datastore is used when restoring to the original VM; otherwise the application randomly selects a datastore attached to the target host.
- **SkipDisks** [<String>] – Sequence number (base 0) of the disks to skip from the savepoint. Disks in savepoint are order by their SCSI controller number and disk number. Multiple numbers are separated by comma.
- **TargetDatastores** [<String>] – Datastores to use on the target host. Multiple values are separated by comma. If only one is supplied, it is used for all disks on the VM. If not set, the original datastores are used when restoring to the original VM; otherwise the application randomly selects a datastore attached to the target host.
- **TargetNetworks** [<String>] – Networks to use on the target host. Multiple values are separated by comma. If only one network name is provided, it is used for all adapters on the VM. If not set, the original networks are used when restoring to the original VM; otherwise the application randomly selects a network attached to the target host.
- **SkipAdapters** [<String>] – Sequence number (base 0) of the adapters to skip in the restore job. Multiple numbers are separated by comma.
- **TransportConfiguration** [<TransportConfiguration>] – The transport configuration options to use for this job.
- **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug,-ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples   -------------------**EXAMPLE 1**-------------------

```
$skipdisks = '0'
$skipnetworks = '0'
 Add-RestorefromManifestJobTemplate -jobname vapirestore_manifest -
username user1 -password Password1 -VmName Test -manifestpath
\\pathtomanifestfile\XYZ-FreeDos_501ec05f-6088-5a14-8ae5-
483bef2e054f\XYZ-FreeDos_20100302_085918_22c5f3ee-9a09-47ae-8a70-
dae9233e5776\XYZ-FreeDos_20100302_085918_F.Manifest.metadata -
targethost $targethost.name -targetDatastores $targetds.name -
PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
Flags $flags -SkipDisks $skipdisks -Skipadapters $skipnetworks -
RunJobNow 1
```

This command creates a restore job template from a manifest file that restores to the specified target host. It also specifies the datastore and network to use on the target host. The first disk (0) in the savepoint and first adapter of the VM are skipped by the restore job.

-------------------**EXAMPLE 2**-------------------

```
Add-RestorefromManifestJobTemplate -jobname vapirestore_manifest -
username user1 -password Password1 -VmName Test -manifestpath
\\pathtomanifestfile\pintsize_20090724093457\pintsize.info -
targethost $targethost.name -targetDatastores $targetds.name -
PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
Flags $flags -RunJobNow 1
```

This command creates a restore job template from a legacy manifest file that restores to the specified target host. It also specifies datastore and network to use on the target host.

-------------------**EXAMPLE 3**-------------------

```
Add-RestorefromManifestJobTemplate -jobname vapirestore_manifest -
username user1 -password Password1 -encryptionpassword encryptPwd -
VmName Test -manifestpath \\pathtomanifestfile\XYZ-
FreeDos_501ec05f-6088-5a14-8ae5-483bef2e054f\XYZ-
FreeDos_20100302_085918_22c5f3ee-9a09-47ae-8a70-dae9233e5776\XYZ-
FreeDos_20100302_085918_F.Manifest.metadata -targethost
$targethost.name -targetDatastores $targetds.name -
PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
Flags $flags -RunJobNow 1
```

 This command creates a restore job template from a encrypted savepoint in the manifest file that restores to the specified target host. It also specifies datastore and network to use on the target host.

Examples – Continued	-------------------EXAMPLE 4-------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
 $targethost = get-inventoryentities -type EsxHost | where{$_.name
-eq 'dev015.dev.eng.int'} $targetds = get-datastore $targethost |
where{$_.Name -eq 'Dynamo-ISCSI'}
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$flags = New-RestoreFlag -ForceOverwrite 1
Add-RestorefromManifestJobTemplate -jobname vapirestore_manifest -
username user1 -password Password1 -VmName Test -manifestpath
\\pathtomanifestfile\XYZ-FreeDos_501ec05f-6088-5a14-8ae5-
483bef2e054f\XYZ-FreeDos_20100302_085918_22c5f3ee-9a09-47ae-8a70-
dae9233e5776\XYZ-FreeDos_20100302_085918_F.Manifest.metadata -
targethost $targethost.name -targetDatastores $targetds.name -
PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
Flags $flags -RunJobNow 1
```

This command creates a restore job template from a manifest file that restores to the specified target host. It also specifies datastore and network to use on the target host.

Related Links	Add-RestoreJobTemplate

Add-RestoreJobTemplate

Description	Adds a restore job template. The user also has the option of running the job immediately after the template is created.
Syntax	

```
Add-RestoreJobTemplate [-JobName] <String>
[[-JobDescription] <String>] [-SavePoint] <SavePoint> [[-
JobSchedule] <Schedule>] [[-Flags] <NoFlags | ForceOverwrite |
ForcePowerOn | UseFiber | PerformNetworkRestoreOnFailure |
AddHypervVirtualMachineToCluster | AllFlags>]
[[-NotificationList] <List`1>] [[-RunJobNow] <Boolean>] [[-VMName]
<String>] [[-TargetHost] <String>] [[-PrimaryDatastore] <String>]
[[-SkipDisks] <List`1>] [[-TargetDatastores] <List`1>]
[[-TargetNetworks] <List`1>] [[-SkipAdapters] <List`1>] [[-
TransportConfiguration] <TransportConfiguration>] [[-
PhysicalMachineIP] <String>] [[-RestoreDiskMapping] <Hashtable>]
[<CommonParameters>]
```

Parameters

- **JobName** <String> – Job name.
- **JobDescription** [<String>] – Job description.
- **SavePoint** <SavePoint> – Savepoint to restore.
- **JobSchedule** [<Schedule>] – Schedule on which the job should run. If not set, the job runs on demand.
- **Flags** [<RestoreFlags>] – Restore flags. If not set, NoFlag. Restore flags: NoFlags, AllFlags, ForceOverwrite, ForcePowerOn. Multiple flags can be separated by '|'
- **NotificationList** [<String>] – Email notification list for the job. Separate multiple emails by ','.
- **RunJobNow** [<Boolean>] – Whether to run the job now. If not set, the job does not run immediately after creation of the template.
- **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used.
- **TargetHost** [<String>] – Fully qualified domain name or IP address of the target host.
- **PrimaryDatastore** [<String>] – Datastore on target host where the restored VM is located. If not set, the original primary datastore is used when restoring to the original VM; otherwise the application randomly selects a datastore attached to the target host.
- **SkipDisks** [<String>] – Sequence number (base 0) of the disks to skip from the savepoint. Disks in savepoint are order by their SCSI controller number and disk number. Multiple numbers are separated by comma.
- **TargetDatastores** [<String>] – Datastores to use on the target host. Multiple values are separated by comma. If only one is supplied, it is used for all disks on the V M. If not set, the original datastores are used when restoring to the original VM; otherwise the application randomly selects a data store attached to the target host.
- **TargetNetworks** [<String>] – Networks to use on the target host. Multiple values are separated by comma. If only one network name is provided, it is used for all adapters on the VM. If not set, the original networks are used when restoring to the original VM; otherwise the application randomly selects a network attached to the target host.
- **SkipAdapters** [<String>] – Sequence number (base 0) of the adapters to skip in the restore job. Multiple numbers are separated by comma.
- **TransportConfiguration** – The transport configuration options to use for this job.
- **PhysicalMachineIP** [<String>] – Physical machine IP address when booted to restore CD to perform the restore to. It must be booted to the restore CD.
- **RestoreDiskMapping** [<Hashtable>] – Which disk in the repository is mapped to corresponding disks in the physical machine.
- **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples

------------------EXAMPLE 1------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type EsxHost | where{$_.name -
eq 'dev015.dev.eng.int'}
$targetds = get-datastore $targethost | where{$_.Name -eq 'Dynamo-
ISCSI'}
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$skipdisks = '0'
Add-RestoreJobTemplate -jobname vapirestore_skipdisk  -savepoint
$mysp -targethost $targethost.name -targetDatastores $targetds.name
-PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
SkipDisks $skipdisks
```

This command creates a restore job template that restore to the specified target host.
The first disk (0) in the savepoint is skipped by the restore job.

------------------EXAMPLE 2------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$spList = get-repositorysavepoint $repo.id
$mysp = ($spList)[1]
Add-RestoreJobTemplate -Jobname vapirestore  -Savepoint $mysp
```

This command creates a restore job template that restore to the original VM in the
savepoint. The job runs on-demand.

------------------EXAMPLE 3------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type EsxHost | where{$_.name -
eq 'dev015.dev.eng.int'}
$targetds = get-datastore $targethost | where{$_.Name -eq 'Dynamo-
ISCSI'}
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
Add-RestoreJobTemplate -jobname vapirestore_totaraget  -savepoint
$mysp -targethost $targethost.name -targetDatastores $targetds.name
-PrimaryDatastore $targetds.name -targetnetworks $targetnw.name
```

This command creates a restore job template that restore to the specified target host. It
also specifies datastore and network to use on the target host.

Examples – Continued	------------------EXAMPLE 4------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type EsxHost | where{$_.name -
eq 'dev015.dev.eng.int'}
$targetds = get-datastore $targethost | where{$_.Name -eq 'Dynamo-
ISCSI'}
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$skipnetworks = '0'
Add-RestoreJobTemplate -jobname vapirestore_skipdisk  -savepoint
$mysp -targethost $targethost.name -targetDatastores $targetds.name
-PrimaryDatastore $targetds.name -targetnetworks $targetnw.name -
Skipadapters $skipnetworks
```

This command creates a restore job template that restores to the specified target host. The first adapter of the VM in the savepoint is skipped by the restore job.

------------------EXAMPLE 5------------------

```
$myInv = Get-InventoryEntities -Type physicalMachine
$myPM = $myInv | where { $_.OriginatingConnection -eq
'10.6.51.178'}n# the physical machine must be booted to the restore
CD for this to work.
$myDiskMap = Get-PhysicalMachineDiskMap $myPM
$myRepo = Get-Repository -Type nfs
$sp = Get-RepositorySavePoint -Id $myRepo.Id
$mySP = ($sp)[0]
$myManifest = Get-SavepointManifest -Id $mySP.id
# Get the mapped physical disks that are the target of the manifest
file map entry when the job is run.
# This cmdlet matches the disks by size.
$myRestoreDiskMap = Get-RestoreDiskMap -ManifestFile $myManifest -
PhysicalDiskMap $myDiskMap
# the user can do something like:
# $myRestoreDiskMap[$myManifest.FileMaps[0]] =
$myDiskMap.PhysicalDisks[1]
# if they want to manually change which disk gets written to by the
restore.
Add-RestoreJobTemplate -JobName Restore178fromNFS -physicalMachine
$myPM -RestoreDiskMap $myRestoreDiskMap -SavePoint $mySP
```

This command creates a restore job template that restores to the specified target physical machine. The restore disk map is a hashtable that has File Map entries from the savepoint as a key and Physical Disk Entities designating which disks should be restored to on the physical machine.

Related Links	Add-BackupJobTemplate
	Add-PhysicalMachine
	Update-RestoreJobTemplate
	Get-Job
	Get-JobTemplate

Add-SFTPRepository

Description	The Add-SftpRepository cmdlet adds a SFTP repository to vRanger. The repository can be later used to store backup files from backup jobs.

Syntax	`Add-SftpRepository [-Server] <string> [[-Port] [<int>]][[-TargetDirectory] [<string>]] [-User] <string> [-Password] <string>[[-TimeOut] [<short>]]`
	`[-Name] <string> [[-Description] [<string>]]`
	`[[-Overwrite] [<Boolean>]] [[-EnableEncryption] [<Boolean>]][[-EncryptionPassword] [<string>]]`
	`[[-EnableCompress] [<Boolean>]] [[-ImportOption] [<string>]]`
	`[<CommonParameters>]`

Parameters	**Server** <String> – The SFTP Server FQDN or IP. **Port** [<int>] – The SFTP Server Port. **TargetDirectory** [<String>] – The SFTP Target Directory. **User** <String> – The SFTP user name. **Password** <String> – The SFTP Password. **TimeOut** [<short>] – Timeout value. **Name** <String> – A friendly name describing this SFTP server. **Description** [<String>] – The Description describing the SFTP server. **Overwrite** [<Boolean>] – Whether to overwrite the existing repository if one exists. **EnableEncryption** [<Boolean>] – Whether to enable Encryption. **EncryptionPassword** <String> – The Encryption Password. **ImportOption** [<String>] – Whether to import the save points as read-write or read only. Possible options 'ReadOnly' and 'ReadWrite'. **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-SftpRepository -Server myserver.dev.corp -TargetDirectory mydir -User myname -Password mypassword -Name myrepo` This command adds a repository on server myserver.dev.corp pointing to directory named mydir using the credentials provided.
Related Links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-CIFSRepository Add-HypervCluster Add-NVSDRepository Add-FTPRepository

Add-VirtualAppforLinuxFLR

Description	The Add-VirtualAppforLinuxFLR command is now obsolete and has been replaced with Add-VirtualAppforLinuxFLRVA.
Related Links	Add-VirtualAppforLinuxFLRVA

Add-VirtualAppforLinuxFLRVA

Description	Adds a virtual appliance to perform Linux FLR.
Syntax	`Add-VirtualAppforLinuxFLRVA [-VirtualApplianceUUID] <string> [<CommonParameters>]`
Parameters	**VirtualApplianceUUID** <String> – Virtual appliance UUID **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-VirtualAppforLinuxFLRVA -VirtualApplianceUUID 564d8ab2-6c7f-5915-0a94-b224b5a71570`
Related Links	None

Add-VirtualCenter

Description	Adds a VirtualCenter to the inventory
Syntax	`Add-VirtualCenter [-DNSorIP] <string> [-Username] <string>[-UserPassword]<string> [[-Port] [<int>]] [<CommonParameters>]`
Parameters	• **DNSorIP** <String> – The DNS or IP for the VirtualCenter. • **Username** <String> – The Username for the VirtualCenter. • **UserPassword** <String> – The User password for the VirtualCenter. • **Port** [<int>] – The SDK Port for the VirtualCenter. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Add-VirtualCenter -DNSorIP 10.6.20.20 -Username administrator -RootPassword p@ssw0rd` This command adds a VirtualCenter to the inventory.
Related Links	Add-DdbRepository Get-InventoryEntities

Disable-Job

Description	Disables enabled jobs.
Syntax	`Disable-Job [-JobTemplateIDs] <GUID[]> [<CommonParameters>]`
Parameters	• **JobTemplateIDs <GUID[]>** – Job Template IDs for which to enable jobs. Disable multiple jobs by passing in comma-separated job template Ids • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	------------------**EXAMPLE 1**-------------------
	
	`Disable-Job -JobTemplateId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'`
	This command disables the job if it is enabled and returns the job template
	entity.
	------------------**EXAMPLE 2**-------------------
	
	`$jobs = Get-JobTemplate | where {$_.IsEnabled -eq 1}`
	`Disable-Job $jobs[0].Id`
	This command disables the first enabled job in the list and returns the job template entity.

Related Links	Enable-Job Remove-VirtualApplianceConfiguration Get-JobTemplate Remove-JobTemplate

Dismount-LinuxVolume

Description	Unmounts VAR disk from the LinuxFLR VA.
Syntax	`Dismount-LinuxVolume [<CommonParameters>]`

Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Dismount-LinuxVolume`
Related Links	Mount-LinuxVolume

Enable-Job

Description	Enables disabled jobs.
Syntax	`Enable-Job [-JobTemplateIDs] <GUID[]> [<CommonParameters>]`
Parameters	• **JobTemplateIDs** <GUID[]> – Job Template IDs for which to enable jobs. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**--------------------
	
	`Enable-Job -JobTemplateId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'`
	
	This command enables the job if it is disabled and returns the job template entity.
	
	-------------------**EXAMPLE 2**--------------------
	
	`$jobs = Get-JobTemplate | where {$_.IsEnabled -eq 0}`
	`Enable-Job $jobs[0].Id`
	
	This command enables the first disabled job in the list and returns the job template entity.

Related Links	Disable-Job Remove-VirtualApplianceConfiguration Get-JobTemplate Remove-JobTemplate

Get-AddressBook

Description	Gets a list of email addresses from the address book.
Syntax	`Get-Addressbook [<CommonParameters>]`
Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-Addressbook`
Related Links	New-EmailAddress

Get-BackupGroupEntity

Description	Returns a list of backup group entities.
Syntax	`Get-BackupGroupEntity [<CommonParameters>]`
Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**------------------- `Get-BackupGroupEntity` This command gets a list of all backup group entities -------------------**EXAMPLE 2**------------------- `Get-BackupGroupEntity	where{$_.Name -eq 'BackupGroupName'}` This command gets a list of backup groups by name
Related Links	None

Get-CatalogSearchData

Description	Returns the list of results for the matched search string, repositories and virtual machines.
Syntax	`Get-CatalogSearchData [-SearchText] <string> [[-RepositoryList] [<string>]] [[-VMList] [-Type] <Entity>] [<CommonParameters>]`
Parameters	• **SearchText** <String> – Search string • **RepositoryList** [<Repository[]>] – Filter by Repositories. If not specified, all repositories are searched. • **VMList** [<Entity[]>] – Filter by entities. If not specified, all entities are searched. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-CatalogSearchdata -SearchText abcd`
	Gets the result set from the catalog database for matched search text. No repository or machine filtering is applied.
	-------------------**EXAMPLE 2**-------------------
	
	`$repo = Get-Repository | where {$_.Name -eq 'Repo1'}`
	`$vmList = Get-InventoryEntities -Type VirtualMachine | {$_.Name -eq 'VM1'}`
	`Get-CatalogSearchdata -SearchText abcd -RepositoryList $repo -VMList $vm`
	Gets the result set from the catalog database for matched search text file red by repository and machine.

Related Links	Remove-Catalog Get-CatalogStatus Set-Cataloging

Get-CatalogStatus

Description	Gets the active status of Catalog for a list of savepoints.
Syntax	`Get-CatalogStatus [-Savepoints] <GUID[]> [<CommonParameters>]`
Parameters	• **Savepoints** <GUID[]> – List of Savepoint IDs • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Get-CatalogStatus -Savepoints 5a62765e-6a04-4df4-af54-46e5a38ad876`
	Gets the active status of the catalog data for the associated savepoint.
Related Links	Get-CatalogSearchData
	Remove-Catalog
	Set-Cataloging

Get-ConfigOption

Description	Returns the config option values for a specified type.
Syntax	`Get-Configoption [-Type] <string> [<CommonParameters>]`
Parameters	• **Type** <String> – Valid types are Database, Email, Resources, Timeout, General, LinuxFLR or Cataloging
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-Configoption -Type Resources`
Related Links	Set-Resources
	Set-Cataloging
	Add-VirtualAppforLinuxFLR

Get-Connection

Description	Returns connection objects to Hosts or VirtualCenters. Without parameters, it returns all configured connections in vRanger. You can also get connections by name or type. Using the two parameters in the same command return an error.
Syntax	`Get-Connection [[-Name] [<string>]] [[-Type] [<string>]] [<CommonParameters>]`
Parameters	• **Name [<String>]** – The name or IP address of the ESX host, vCenter, Hyper-V Host, or Hyper-V Cluster.
	• **Type [<String>]** – Connection type: EsxHost, VirtualCenter, HypervHost, HypervCluster, or All
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	--------------------EXAMPLE 1--------------------
	`Get-Connection -type HypervHost`
	This command gets a list of all Hyper-V host connections.
	--------------------EXAMPLE 2--------------------
	`Get-Connection`
	This command gets a list of all vRanger loaded connections.
	--------------------EXAMPLE 3--------------------
	`Get-Connection -Name myhost.dev.corp`
	This command returns a connection to the specified host
	--------------------EXAMPLE 4--------------------
	`Get-Connection -type EsxHost`
	This command gets a list of all ESX host connections.
Related Links	None

Get-CurrentTemplateVersionID

Description	Returns the current template version ID of a job template based on a historic template version ID. A new template version ID is assigned to the job template when it is modified.
Syntax	`Get-CurrentTemplateVersionID [-TemplateVersionID] [<string>] [<CommonParameters>]`
Parameters	• **TemplateVersionID** <GUID> – Template version ID.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-CurrentTemplateVersionID 6baedeed-227a-4350-ba3f-470f9e90dc07`
	This command returns the current template version id of the job template referenced by the input template version id.
Related Links	Get-Job
	Get-JobTemplate

Get-Datastore

Description	The Get-Datastore cmdlet returns all datastores associated with the specified host.
Syntax	`Get-Datastore [-Host] <HostEntity> [<CommonParameters>]`
Parameters	• **Host** <HostEntity> – The host object.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-Datastore -Host $host`
	This command returns a list of datastores associated with the host
Related Links	Get-MonitorLog

Get-GlobalTransportFailover

Description	Gets the default Transport Failover order for replication jobs.
Syntax	`Get-GlobalTransportFailover [<CommonParameters>]`
Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-GlobalTransportFailover` This command gets the default Transport Failover order for replication jobs.
Related Links	Update-GlobalTransportFailover Get-TransportFailover

Get-InventoryEntities

Description	Returns a list of inventory entities based on the input repository type. The command returns all inventory entities if no type is provided. You can also get inventory entities by type. Note that there's a performance delay if all types are requested or no type is specified.
Syntax	`Get-InventoryEntities [[-Type] <String>] [[-Block] <Boolean>] [[-Thin] <Boolean>][<CommonParameters>]`
Parameters	• **Type [<String>]** – Inventory entity type: All, ClusterComputeResource, ComputeResource, DataCenter, EsxHost, Folder, HypervCluster, HypervHost, HypervVirtualMachine, ResourcePool, VirtualCenter, VirtualMachine, VirtualApp, PhysicalMachine • **Block [<Boolean>]** – If set to true, the call blocks if inventory is being refreshed until it's done. Possible timeout may occur on large inventory sets. • **Thin [<Boolean>]** – If set to true, the call returns thin inventory objects. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-InventoryEntities - Type HypervHost`
	
	This command gets a list of all inventory entities
	-------------------**EXAMPLE 2**-------------------
	
	`Get-InventoryEntities -Type EsxHost`
	This command gets a list of all loaded ESX host entities.
	-------------------**EXAMPLE 3**-------------------
	
	`Get-InventoryEntities`
	This command gets a list of all inventory entities

Related Links	

Get-IsInventoryRefreshing

Description	The Get-IsInventoryRefreshing cmdlet returns "true" or "false" based on whether the inventory is refreshing at the time the cmdlet is run.
Syntax	

Parameters	• **Type** [<String>] – Inventory entity type to check: All, VMware, Hyper-V
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-IsInventoryRefreshing -Type Hyperv`
	
	`This command returns true if Hyper-V is refreshing.`
	-------------------**EXAMPLE 2**-------------------
	
	`Get-IsInventoryRefreshing`
	
	`This command returns true if any inventory entity is refreshing.`
	-------------------**EXAMPLE 3**-------------------
	
	`Get-IsInventoryRefreshing -Type VMware`
	
	`This command returns true if VMware is refreshing.`

Related Links	

Get-Job

Description	The Get-Job cmdlet returns jobs. Without parameters, Get-Job returns all past and current jobs. You can also get jobs by JobState within a time span or JobTemplateIDs. If JobTemplateIDs are specified, no other parameters can be used.
Syntax	`Get-Job [[-StartTime] [<string>]] [[-EndTime] [<string>]] [[-JobState] [<string>]] [[-JobTemplateIDs] [<string>]] [<CommonParameters>]`
Parameters	• **EndTime [<String>]** – End time of time period.
	• **StartTime [<String>]** – Start time of time period
	• **JobState [<String>]** – Job state: Queued, Running, Completed or All.
	• **JobTemplateIDs [<String>]** – Job template IDs for which to retrieve jobs. Separate multiple IDs by ','.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-Job`
	This command returns jobs in all states
	-------------------**EXAMPLE 2**-------------------
	
	`Get-Job -JobState Running`
	This command returns all running jobs
	-------------------**EXAMPLE 3**-------------------
	
	`Get-Job -starttime 2009-03-30 -endtime 2009-03-31 -JobState Completed`
	This command returns completed jobs that started during the specified time span
	-------------------**EXAMPLE 4**-------------------
	
	`Get-Job d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3`
	This command returns jobs based on specified template ID.

Related Links	Get-JobTemplate

Get-JobTemplate

Description	The Get-JobTemplate cmdlet returns a list of all jobs, on demand and scheduled (job templates). Without parameters, the command returns all job templates. You can also get job templates by Type or Inventory IDs.
Syntax	`Get-JobTemplate [[-Type] <String>] [[-InventoryIds] <String[]>] [[-Thin] <Boolean>] [[-ID] <String>] [-VersionID] <String>] [<CommonParameters>]`
Parameters	• **Type [<String>]** – Specifies the type of job template to retrieve: Backup, Restore or All. • **InventoryIds [<String>]** – The inventory nodes' virtualization object ref IDs. Separate multiple IDs by ','. • **Thin [<Boolean>]** – Whether to return 'thin' version of job templates. Returning thin objects can speed up commands and avoid performance related errors. By default, full objects are returned. • **ID [<String>]** – Job Template ID • **VersionID [<String>]** – Job Template version ID • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-JobTemplate`
	This command gets a list of all job templates
	-------------------**EXAMPLE 2**-------------------
	
	`Get-JobTemplate -type Backup`
	This command gets a list of all backup jobs.
	-------------------**EXAMPLE 3**-------------------
	
	`Get-JobTemplate -InventoryIds d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-24 12990005f3 -Type Backup`
	This command returns backup job templates in the specified inventories

Related Links	Get-Job

Get-MonitorLog

Description	
Syntax	`Get-MonitorLog [[-EventId] <Int32>] [[-StartDateTime] <DateTime>] [[-EndDateTime] <DateTime>] [-Verbose] [-Debug] [-ErrorAction <ActionPreference>] [-WarningAction <ActionPreference>] [-ErrorVariable <String>] [-WarningVariable <String>] [-OutVariable <String>] [-OutBuffer <Int32>] [-WhatIf] [-Confirm]`
Parameters	
Examples	
Related Links	

Get-Network

Description	The Get-Network cmdlet returns all networks associated with the host specified.
Syntax	`Get-Network [-Host] <HostEntity> [<CommonParameters>]`
Parameters	• **Host <HostEntity>** – The host object. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-Network -Host $host` This command returns a list of networks associated with the host
Related Links	Get-Datastore

Get-PhysicalMachineDiskMap

Description	Creates a Physical Machine disk map to be used in AddRestoreJobTemplate
Syntax	`Get-PhysicalMachineDiskMap [-RestoreCDIP] <String> [<CommonParameters>]`
Parameters	• **RestoreCDIP** <String> – The physcial machine IP address as shown when it is booted to the restore CD • **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`GetPhysicalMachineDiskMap -RestoreCDIP $IP` Creates a Physical Machine disk map to be used in AddRestoreJobTemplate.
Related Links	Get-InventoryEntities Add-RestoreJobTemplate

Get-Repository

Description	Returns repositories configured in vRanger. Without parameters, "Get-Repository" gets all repositories configured in vRanger. You can also get repositories by Name or Type. Using the two parameter in the same command returns an error.
Syntax	`Get-Repository [[-Id] [<string>]] [[-Type] <string>] [<CommonParameters>]`
Parameters	• **Type <String>** – Specifies the type of repository to retrieve: CIFS, SFTP, FTP, NFS, NVSD, DDB, or ALL. • **Id <String>** – Repository ID in GUID. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`Get-Repository`
	
	This command gets a list of all vRanger repositories
	
	-------------------**EXAMPLE 2**-------------------
	
	`Get-Repository -id 034538e5-58a7-4d28-89c9-3249b620064a`
	
	This command a repository with the specified ID
	
	-------------------**EXAMPLE 3**-------------------
	
	`Get-Repository -type CIFS`
	
	This command gets a list of all CIFS repositories.

Related Links	Add-CIFSRepository Add-DdbRepository Add-FTPRepository Add-HypervCluster Add-SFTPRepository Remove-Repository

Get-RepositoryJob

Description	The Get-RepositoryJob cmdlet returns a list of associated jobs based on the input repository ID.
Syntax	`Get-RepositoryJob [-RepositoryId] <string> [<CommonParameters>]`
Parameters	• **RepositoryId <String>** – Repository ID in GUID. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-RepositoryJob -RepositoryId 034538e5-58a7-4d28-89c9-3249b620064a` This command returns a list of jobs associated with the repository with the specified ID
Related Links	Get-Job Get-JobTemplate

Get-RepositorySavePoint

Description	The Get-RepositorySavePoint cmdlet returns a list of savepoints in the repository. All savepoints in the repository are returned if no time parameter is specified.
Syntax	`Get-RepositorySavePoint [-Id] <string> [[-StartTime] [<string>]] [[-EndTime] [<string>]] [<CommonParameters>]`
Parameters	• **Id** <String> – The repository Id to Retrieve information for. • **StartTime** [<String>] – Start time of timespan during which the savepoints are complete • **EndTime** [<String>] – End time of timespan during which the savepoints are completed • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Get-RepositorySavePoint -Id 034538e5-58a7-4d28-89c9-3249b620064a`
	This command returns a list of savepoints associated with the repository with the specified ID
Related Links	None

Get-RestoreDiskMap

Description	Creates a Hashtable that represents a matching of file map entries from the savepoint with the disks from a physical machine (a RestoreDiskMap).
Syntax	`Get-RestoreDiskMap [-PhysicalDiskMap] <PhysicalDiskMap> [-ManifestFile] <ManifestFile> [<CommonParameters>]`
Parameters	• **PhysicalDiskMap** <PhysicalDiskMap> – The physical machine disk map
	• **ManifestFile** <ManifestFile> – The manifest from the savepoint that contains file map entries for the disks backed up.
	• **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-RestoreDiskMap -PhysicalDiskMap $physicaldiskMap -ManifestFile $manifestFile`
	This command matches the file map entries to the disks in a physical machine by size and returns a hash map that the user can manipulate by code
Related Links	Add-Backup
	Get-InventoryEntities

Get-SavepointDisk

Description	Returns the VM disks that the savepoint contains.
Syntax	`Get-SavepointDisk [-Id] <String> [<CommonParameters>]`
Parameters	• **Id <String>** – SavePoint ID in GUID format
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-SavePointDisk 91d27db5-4562-4456-918d-b0dc587e0e7f`
	This command returns all disks that has been backed up into the savepoint specified
Related Links	Get-VMDKVolume

Get-SavepointManifest

Description	Get-SavePointManifest cmdlet returns manifest of the savepoint
Syntax	`Get-SavePointManifest [-ID] <ID> [<CommonParameters>]`
Parameters	• **ID <ID>** – The SavePoint ID.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	Get-SavePointManifest $id
	This command returns savepoint manifest object
Related Links	Get-RepositorySavePoint
	Get-VMDKVolume

Get-Savepoints

Description	Returns a list of savepoints for a specified BackupJobTemplateID
Syntax	Get-SavePoints [-BackupJobTemplateID] <GUID> [<CommonParameters>]
Parameters	• **BackupJobTemplateID <GUID>** – The backup job template ID to retrieve savepoints for.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	Get-SavePoints -BackupJobTemplateID 034538e2-58a7-4d28-89c9-3249b620064a
	This command returns a list of savepoints associated with the backup template with the specified ID.
Related Links	Get-Job
	Get-JobTemplate

Get-TransportFailover

Description	Gets the Transport Failover order for replication jobs by TemplateVersion ID.
Syntax	Get-TransportFailover [-TemplateVersionID] <string> [<CommonParameters>]
Parameters	• **TemplateVersionID** <String> – Template Version Id for job.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	$jobtemplates = Get-JobTemplate -Type Replication $template = $jobtemplates[0] Get-TransportFailover -TemplateVersionID $template.TemplateVersionId
	Gets the Transport Failover order for replication jobs by TemplateVersion ID.
Related Links	Get-GlobalTransportFailover
	Update-GlobalTransportFailover

Get-VirtualApplianceConfig

Description	Gets the VirtualApplianceConfig configuration from vRanger.
Syntax	Get-VirtualApplianceConfig [[-VAName] [<string>]] [[-HostName] [<string>]] [<CommonParameters>]

Parameters	• **VAName** [<String>] – The virtual appliance Name.
	• **HostName** [<String>] – The Host Name where the virtual appliance resides.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	-------------------**EXAMPLE 1**-------------------
	`$va = Get-VirtualApplianceConfig`
	This command gets all the virtual appliance entities.
	-------------------**EXAMPLE 2**-------------------
	`$va = Get-VirtualApplianceConfig -VAName ABC-Vzva-1.1.1`
	This command gets the virtual appliance entity for the VA specified.
	-------------------**EXAMPLE 3**-------------------
	`$va = Get-VirtualApplianceConfig -HostName xyz.dev.eng.int`
	This command gets the virtual appliance entities for the Host specified.
Related Links	New-VirtualAppliance
	Update-VirtualAppliance
	Remove-VirtualAppliance

Get-VirtualApplianceDeploymentStatus

Description	Gets the status of the re-configuration of VA in progress
Syntax	`Get-VirtualApplianceDeploymentStatus [-DeploymentId] <GUID> [<CommonParameters>]`
Parameters	• **ID** <GUID> – Deployment Id for VA
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-VirtualApplianceReconfigStatus -Id $deploymentID`
	This command gets the status of the VA re-configuration.
Related Links	Update-VirtualAppliance
	New-VirtualAppliance

Get-VirtualApplianceReconfigStatus

Description	Gets the status of the deployment of VA in progress
Syntax	`Get-VirtualApplianceDeploymentStatus [-DeploymentId] <GUID> [<CommonParameters>]`
Parameters	• **DeploymentId** <GUID> – Deployment Id for VA
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-VirtualApplianceDeploymentStatus -Id $deploymentID`
	This command gets the status of the VA deployment in progress.
Related Links	Update-VirtualAppliance
	New-VirtualAppliance

Get-VirtualMachinesUnderInventory

Description	Gets the child VM entities of the entity (host, cluster, etc) provided.
Syntax	`Get-VirtualMachinesUnderInventory [[-InventoryEntityMoRefId] [<string>]] [<CommonParameters>]`
Parameters	• **InventoryEntityMoRefId** [<String>] – The MoRef ID of the parent inventory entity. • **Thin [<Boolean>}** – If set to true, the call will return thin inventory objects. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	-----------------EXAMPLE------------------ `Get-VirtualMachinesUnderInventoryEntity -InventoryEntityMoRefId host-16641` The managed object reference ID of an inventory entity node such as a ESX host, or resource folder.
Related Links	Add-BackupGroupEntity Get-BackupGroupEntity Get-InventoryEntities

Get-VmDisk

Description	The Get-VmDisk cmdlet returns a list of VmDisks for the VM list in the same order of the input.
Syntax	`Get-VmDisk [-VMList] <VMEntity[]>] [<CommonParameters>]`
Parameters	• **VMList <VMEntity[]>** – A list of VM entities to retrieve disk information for. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Get-VmDisk -VMList $mylist` This command gets a list of VMDisks objects for the input VM list
Related Links	None

Get-VMDKVolume

Description	Returns all volumes on a certain VM disk.
Syntax	`Get-VMDKVolume [-VmdkID] <String> [-Repository] <Repository> [-RepositoryPassword] <String>[SavepointManifest] <ManifestFile> [<CommonParameters>]`
Parameters	• **VmdkID <String>** – UUID of the VMDK • **Repository <Repository>** – Repository in which the SavePoint resides in • **RepositoryPassword <String>** – Repository password • **SavepointManifest** <ManifestFile> – The savepoint manifest object • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Get-VMDKVolume -vmdkid $diskid -repository $repo -`
	`repositorypassword Password -savepointmanifest $manifest`
	This command returns all volumes of the specified VMDK

Related Links	None

Install-VirtualAppliance

Description	Deploys a new virtual appliance for vRanger.

Syntax	`Install-VirtualAppliance [-VAName] <String> [-SelectedHost]`
	`<HostEntity> [-OVFLocation] <String> [-DownloadOVF] <Boolean> [[-`
	`IPAddressOverride] <String>] [[-ClusterDefault] <Boolean>] [-`
	`CPUAllocation] <Int32> [-MemorySizeInMB] <Int64> [-VADatastore]`
	`<Datastore> [-NetworkEntity] <Network> [[-ScratchDiskDatastore]`
	`<Datastore>] [[-ScratchDiskSizeInKB] <Int64>] [[-NewPassword]`
	`<String>] [[-UseVAForLFLR] <Boolean>] [[-PowerOnVA] <Boolean>] [[-`
	`VADeployStatus] <Boolean>] [[-UseDHCP] <Boolean>] [[-NetMask]`
	`<String>] [[-StaticIP] <String>] [[-GateWay] < String>] [[-`
	`UseStaticDNS] <Boolean>] [[-PrimaryDNS] <String>] [[-SecondaryDNS]`
	`<String>] [[-Domain] <String>] [[-Search] <String>]`
	`[<CommonParameters>]`

Parameters	• **VAName <String>** – The virtual appliance name

Parameters

- **VAName <String>** – The virtual appliance name
- **SelectedHost <HostEntity>** – The host to associate the virtual appliance with
- **OVFLocation <String>** – The path to the local OVF
- **DownloadOVF** <Boolean> – Should the OVF be downloaded?
- **IPAddressOverride** [<String>] – IP Address to use for the virtual appliance instead of the one retrieved from the VirtualCenter API.
- **ClusterDefault** [<Boolean>] – If the selected host belongs to a cluster, share the virtual appliance among all hosts in the cluster
- **CPUAllocation** <Int> – Allocate CPU
- **MemorySizeInMB** <Int> – Assign memory in MB
- **VADatastore** <Datastore> – Select a datastore for the virtual appliance.
- **NetworkEntity** <Network> – Select a network for the virtual appliance.
- **ScratchDiskDatastore** [<Datastore>] – Select a datastore for the ScratchDisk.
- **ScratchDiskSizeInKB** [<Int>] – Assign ScratchDisk in KB.
- **NewPassword** [<String>] – New password for virtualAppliance, if this is not mentioned, Original password (vzroot1) is used
- **UseVAForLFLR** [<Boolean>] – Use virtual appliance for Linux FLR.
- **PowerOnVA** [<Boolean>] – Power On virtual appliance after deployment complete
- **VADeployStatus** [<Boolean>] – Retrieves VA deployment status if set to true.Defaults to false if the parameter is not passed in.
- **UseDHCP** [<Boolean>] – Specify use of DHCP to set IP address. If false, StaticIP, NetMask, Gateway and PrimaryDNS must be specified.
- **NetMask** [<String>] – Net mask for use with the VA's static IP address. Required if UseDHCP is false.
- **StaticIP** [<String>] – Static IP Address for the VA. Required if UseStaticIP is true.
- **GateWay** [<String>] – Optional gateway address for VA.
- **UseStaticDNS** [<Boolean>] – Specify use of a static DNS setup. If true, PrimaryDNS must be specified.
- **PrimaryDNS** [<String>] – Mandatory address of primary DNS server for VA. Required if UseStaticDNS is true.
- **SecondaryDNS** [<String>] – Optional address of secondary DNS server for VA.
- **Domain** [<String>] – Optional Local domain name for DNS resolver on VA. See Linux man page for resolv.conf for more info.
- **Search** [<String>] – Optional Search list for host-name lookup for DNS on VA. See Linux man page for resolv.conf for more info.
- **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$myhost = Get-InventoryEntities -Type esxhost | where{$_.Name -eq 'dev-pool001.dev.eng.int'}`
	`$ds=Get-Datastore $myhost | where{$_.Name -like 'Local*'}`
	`$nw=Get-Network -Host $myhost Install-VirtualAppliance -VAName kk-TestVAStaticIP -SelectedHost $myhost -OVFLocation`
	`'X:\Development\Vzva\vzva-1.8.0.ova' -DownloadOVF 0 -CPUAllocation`
	`1 -MemorySizeInMB 512 -VADatastore $ds -NetworkEntity $nw -`
	`PowerOnVA 1 -VADeployStatus 1 -UseDHCP 0 -staticIP 10.6.51.159 -`
	`NetMask 255.255.254.0 -Gateway 10.6.50.2 -UseStaticDNS 1 -`
	`PrimaryDNS 10.6.30.10`
	
	This command deploys a new VirtualAppliance configuration to vRanger using a locally saved OVF and retrieves the status of the currently deploying VA in progress. Default for VADeployStatus is false

Related Links	Get-VirtualApplianceConfig Remove-VirtualAppliance Update-VirtualAppliance

Mount-LinuxVolume

Description	Mounts a volume from the specified manifest FileMaps.
Syntax	`Mount-LinuxVolume [-ManifestFileMaps] <FileMapEntryLine[]>[-Repository] <Repository>` `[-RepositoryPassword] <String> [<CommonParameters>]`
Parameters	• **ManifestFileMaps** <FileMapEntryLine[]> – The savepoint manifest file map • **Repository** <Repository> – The repository containing the VAR • **RepositoryPassword** <String> – The password to the repository • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`$path = Mount-LinuxVolume -repositorypassword xxxxxxxx -` `manifestfilemap [filemap] -repository [repo]`
Related Links	Dismount-LinuxVolume

New-BackupFlag

Description	The New-BackupFlag cmdlet creates a new backup flag for later use.

Syntax	`New-BackupFlag [[-BackupPoweredOnOnly] <Boolean>]`
	`[[-CheckDestinationFreeSpace] <Boolean>]`
	`[[-UseCompression] <Boolean>] [[-UpdateNotes] <Boolean>] [[-`
	`EnableGuestQuiescing] <Boolean>]`
	`[[-EnableABM] <Boolean>] [[PerformNetworkBackupOnFailure]`
	`<Boolean>]`
	`[[-UseFiber] <Boolean>] [[-UseCataloging] <Boolean>] [[-LanFree]`
	`<Boolean>] [[-AllFlags] <Boolean>] [<CommonParameters>]`

Parameters	• **BackupPoweredOnOnly** [<Boolean>] – Whether to back up only virtual machines that are powered on.
	• **CheckDestinationFreeSpace** [<Boolean>] – Whether to check repository free space before backup.
	• **UseCompression** [<Boolean>] – Whether to use compression for backup.
	• **UpdateNotes** [<Boolean>] – Whether to update notes on VM after backup.
	• **EnableGuestQuiescing** [<Boolean>] – Whether guest quiescing is enabled.
	• **EnableABM** [<Boolean>] – Whether Active Block Mapping is enabled.
	• **PerformNetworkBackupOnFailure** [<Boolean>] – [Deprecated] Formerly, whether to perform the backup using the network if there is a problem during the fiber or iSCSI backup.
	• **UseFiber** [<Boolean>] – [Deprecated] Formerly, whether to use fiber or iSCSI to perform this backup. Currently, whether to perform a Lan-Free backup. Use **LanFree** instead.
	• **UseCataloging** [<Boolean>] – Whether to use cataloging to perform this backup.
	• **EnableVSSApplicationLevelQuiescing** [<Boolean>] – Whether to enable VSS Application Level Quiescing for this backup
	• **EnableVSSLogTruncation** [<Boolean>] – Whether to enable VSS log truncation for this backup
	• **LanFree** [<Boolean>] – Whether to perform a Lan-Free backup.
	• **AllFlags** [<Boolean>] – Whether to use all available flags for the backup.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	-------------------**EXAMPLE 1**-------------------
	`New-BackupFlag`
	This command returns a new BackupFlags object with the value of NoFlags.
	-------------------**EXAMPLE 2**-------------------
	`New-BackupFlag -EnableGuestQuiescing 1 -EnableABM 1`
	This command returns a new BackupFlags object with the values of EnableGuestQuiescing and EnableABM.
	-------------------**EXAMPLE 3**-------------------
	`New-BackupFlag -AllFlags 1`
	This command returns a new BackupFlags object with all flags set to on.
Related Links	Add-BackupJobTemplate

New-BackupGroupMember

Description	The New-BackupGroupMember cmdlet is used to add Members to a Backup Group created using Add-BackupGroupEntity cmdlet or UI. User can add any Inventory item as a member.
Syntax	`New-BackupGroupMember [-BackupGroupId] <GUID> [-MemberEntity] <Entity> [<CommonParameters>]`
Parameters	• **BackupGroupId** <GUID> – The BackupGroup Id
	• **MemberEntity** <Entity> – Inventory entity to add to Backup Group
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$inventoryEntity = Get-InventoryEntity -type virtualmachine | where {$_.Name -eq 'VMNAMEHERE'}`
	
	`New-BackupGroupMember 8c40f593-6608-45ef-b2c9-a2594aac2338 $inventoryEntity`

Related Links	None

New-Daily Schedule

Description	Returns a Schedule object. The schedule provides a template daily schedule that can be further customized before use.
Syntax	`New-DailySchedule [-WeekdayOnly] <Boolean> [-StartTime] <string> [<CommonParaeters>]`
Parameters	• **WeekdayOnly <Boolean>** – Whether to run the job only on weekdays. • **StartTime <String>** – Job start time. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`New-DailySchedule -WeekdayOnly 1 -StartTime "2009-03-31 12:00:00AM"` This command returns a daily job schdule that runs weekdays at midnight.
Related Links	None

New-EmailAddress

Description	The New-EmailAddress cmdlet is used to add a new email address to the address book.
Syntax	`New-EmailAddress [-Name] <string> [-Email] <string> [<CommonParameters>]`
Parameters	• **Name** <String> – Name of Person • **Email** <String> – Email address • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`New-EmailAddress -Name Employee1 -Email firstname.lastname@company.com`
Related Links	Get-AddressBook

New-IntervalSchedule

Description	The New-IntervalSchedule cmdlet returns a interval Schedule object.
Syntax	`New-IntervalSchedule [[-IntervalDays] <int>] [[-IntervalHours] [<int>]] [-IntervalMinutes] <int> [-StartTime] <string> [[-RangeEndDate] <string>] [[-RangeByOccurrence] <int>] [<CommonParameters>]`

Parameters	• **IntervalDays** <int> – Number of days to use for interval schedule. Default is 0.
	• **IntervalHours** <int> – Number of hours to use for interval schedule. Default is 0.
	• **IntervalMinutes** <int> – Number of minutes to use for interval schedule. Default is 0.
	• **StartTime** <String> – Job start time.
	• **RangeEndDate** <String> – Range end date.
	• **RangeByOccurrence** <int> – The maximum number of times a job is executed.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	------------------**EXAMPLE 1**------------------
	`New-IntervalSchedule -StartTime 2009-03-31 -IntervalHours 1`
	This command returns a interval job schdule that runs every hour.
	------------------**EXAMPLE 2**------------------
	`New-IntervalSchedule -StartTime 2009-03-31 -IntervalMinutes 30 -` `RangeEndDate 2010-12-31`
	This command returns a interval job schdule that runs every 30 minutes till 2010-12-31
	------------------**EXAMPLE 3**------------------
	`New-IntervalSchedule -StartTime 2009-03-31 -IntervalHours 1 -` `intervalMinutes 30 -RangeByOccurrence10`
	This command returns an interval job schdule that runs every hour and half; the job is executed 10 times at this interval.
Related Links	New-Daily Schedule
	New-WeeklySchedule
	New-MonthlySchedule
	New-YearlySchedule

New-MonthlySchedule

Description	The New-MonthlySchedule cmdlet returns a Schedule object. The schedule provides a template monthly schedule that can be further customized before use.
Syntax	`New-MonthlySchedule [-ExecutionDayOfMonth] <int> [-StartTime]` `<string> [<CommonParameters>]`
Parameters	• **ExecutionDayOfMonth <int>** – Day of month on which the job runs (1-31).
	• **StartTime <String>** – Job start time.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`New-MonthlySchedule -ExecutionDayOfMonth 1 -StartTime "2009-03-31` `12:00:00AM"`
	This command returns weekly job schdule that runs on the first of every month at midnight
Related Links	New-Daily Schedule
	New-WeeklySchedule
	New-YearlySchedule
	New-IntervalSchedule

New-ReplicationFlag

Description	The New-ReplicationFlag cmdlet helps the user create a ReplicationFlags Enum object. Without parameters, the command returns NoFlags. You can pass in boolean parameters to specify which flags you want on.
Syntax	`New-ReplicationFlag [[-CheckDestinationFreeSpace] [<Boolean>]] [[-UpdateNotes] [<Boolean>]] [[-EnableGuestQuiescing] [<Boolean>]] [[-EnableABM] [<Boolean>]] [[-AllFlags] [<Boolean>]] [<CommonParameters>]`
Parameters	• **CheckDestinationFreeSpace** [<Boolean>] – Whether to check repository free space before replication. • **UpdateNotes** [<Boolean>] – Whether to update notes on VM after replication. • **EnableGuestQuiescing** [<Boolean>] – Whether guest quiescing is enabled. • **EnableABM** [<Boolean>] – Whether active block mapping is enabled. • **AllFlags** [<Boolean>] – Whether to use all available flags for the replication. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**-------------------
	
	`NewReplicationFlag`
	This command returns a new ReplicationFlags object with the value of NoFlags
	-------------------**EXAMPLE 2**-------------------
	
	`New-ReplicationFlag -EnableGuestQuiescing 1 -EnableABM 1`
	This command returns a new ReplicationFlags object with the values of EnableGuestQuiescing and EnableABM
	-------------------**EXAMPLE 3**-------------------
	
	`New-ReplicationFlag -AllFlags 1`
	This command returns a new ReplicationFlags object with all flags set to on

Related Links	Add-PhysicalMachine

New-RestoreFlag

Description	The New-RestoreFlag cmdlet helps the user create a RestoreFlags Enum object. Without parameters, the command returns NoFlags. You can pass in boolean parameters to specify which flags you want on.
Syntax	`New-RestoreFlag [[-ForceOverwrite] <Boolean>] [[-ForcePowerOn] <Boolean>] [[-UseFiber] <Boolean>] [[-PerformNetworkRestoreOnFailure] <Boolean>] [[-LanFree] <Boolean>] [[-AllFlags] <Boolean>] [<CommonParameters>]`

Parameters	• **ForceOverwrite [<Boolean>]** – Whether to force overwrite of existing VM with same name.
	• **ForcePowerOn [<Boolean>]** – Whether to power on the VM for restore.
	• **UseFiber** [<Boolean>] – [Deprecated] Formerly, whether to use fiber or iSCSI to perform this restore. Currently, whether to perform a Lan-free Restore. Use LanFree instead.
	• **PerformNetworkRestoreOnFailure** [<Boolean>] – [Deprecated] Formerly, whether to perform a network restore on a Lan-free restore failure.
	• **LanFree** [<Boolean>] – Whether to perform a Lan-free Restore.
	• **AllFlags [<Boolean>]** – Whether to use all available flags for the restore.
	• **AddHypervVirtualMachineToCluster** [<Boolean>} – Whether to add the restored Hyper-V virtual machine to the cluster.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	------------------**EXAMPLE 1**------------------
	
	`New-RestoreFlag`
	
	This command runs RestoreFlags with the value of NoFlags
	
	------------------**EXAMPLE 2**------------------
	
	`New-RestoreFlag -ForcePowerOn 1 -UseCompression 0`
	
	This command runs RestoreFlags with the values of ForcePowerOn
	
	------------------**EXAMPLE 3**------------------
	
	`New-RestoreFlag -AllFlags 1`
	
	This command runs RestoreFlags with all flags set to on.

Related Links	None

New-SMTPServer

Description	The New-SMTPServer cmdlet is used to add a new SMTP Server.
Syntax	`New-SMTPServer [-SMTPServer] <string> [-Port] <int> [-FromEmail] <string> [<CommonParameters>]`
Parameters	• **SMTPServer** <String> – SMTP Server address
	• **Port** <int> – Port No for SMTP Server.
	• **FromEmail** <String> – Email address for sent emails.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`New-SMTPServer -SMTPServer company.smtp.com -Port 22 -FromEmail support@company.com`
Related Links	None

New-TransportConfiguration

Description	Helps the user create a TransportConfiguration object. Without parameters, the command returns the default Automatic configuration. You can pass in boolean parameters to specify which settings you want on.
Syntax	`New-TransportConfiguration [-AutoSelectTransport] <Boolean> [-RunLocal] <Boolean>[-UseAdvancedTransport] <Boolean> [-UseLanTransport] <Boolean>` `[[-UseSsLforLan] <Boolean>] [<CommonParameters>]`
Parameters	• **AutoSelectTransport** <Boolean> – The service determines which transport type to use at run-time. • **RunLocal** <Boolean> – Whether to run the job on the local machine vs the VA. • **UseAdvancedTransport** <Boolean> – Whether the job should attempt to use an advanced transport (SAN/HotAdd) first. • **UseLanTransport** <Boolean> – Whether the job should attempt to use a network transport. • **UseSsLforLan** [<Boolean>] – Whether the job should use an encrypted connection when attempting a network transport. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	------------------**EXAMPLE 1**------------------
	`New-BackupFlag -RunLocal 1 -UseAdvancedTransport 1 UseLANTransport 0`
	This command returns a new transport configuration object with the values of RunLocal, UseAdvancedTransport.
	------------------**EXAMPLE 2**------------------
	`New-TransportConfiguration`
	This command returns a new transport configuration object with the configuration set to Automatic.

Related Links	Add-BackupJobTemplate Add-RestoreJobTemplate

New-VirtualAppliance

Description	Configures a virtual appliance for vRanger.
Syntax	`New-VirtualAppliance [-VAEntity] <Entity>` `[-RootPassword] <String> [[-IPAddressOverride] <String>] [[-ClusterDefault] <Boolean>] [<CommonParameters>]`
Parameters	• **VAEntity** <Entity> – The VM Entity where the virtual appliance is deployed. • **RootPassword** <String> – Root Password for the virtual appliance. • **IPAddressOverride** <String> – New IP Address for the VA • **ClusterDefault** [<Boolean>] – Share the virtual appliance among all hosts in the cluster. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

| Examples | ```
$va = Get-InventoryEntities -Type VirtualMachine | where {$_.Name -
eq 'ABC-Vzva-1.1.1'}
New-VirtualAppliance -VAEntity $va -IPAddressOverride 10.6.50.19
-RootPassword vzroot1
```<br><br>This command adds a virtual appliance configuration to vRanger. |
| Related Links | Update-VirtualAppliance
Remove-VirtualAppliance
Get-VirtualApplianceConfig |

New-WeeklySchedule

Description	The New-WeeklySchedule cmdlet returns a Schedule object. The schedule provides a template weekly schedule that can be further customized before use.
| Syntax | ```
New-WeeklySchedule [-ExecutionDays] <string> [-StartTime] <string>
[<CommonParameters>]
``` |
| Parameters | • **ExecutionDays <String>** – Days of the week on which the job runs: Monday, Tuesday, and so on. Operate multiple days by ','.
• **StartTime <String>** – Job start time.
• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters". |
| Examples | -------------------**EXAMPLE 1**-------------------<br><br>```
NewWeeklySchedule -ExecutionDays "Monday" -StartTime "2009-03-31
12:00:00AM"
```<br><br>This command returns weekly job schdule that runs on Mondays at midnight<br><br>-------------------**EXAMPLE 2**-------------------<br><br>```
NewWeeklySchedule -ExecutionDays "Monday,Tuesday" -StartTime "2009-
03-31 12:00:00AM"
```<br><br>This command returns weekly job schdule that runs on Mondays and Tuesdays at midnight |
| Related Links | None |

New-YearlySchedule

Description	The NewYearlySchedule cmdlet returns a Schedule object. The schedule provides a template yearly schedule that can be further customized before use.
| Syntax | ```
NewYearlySchedule [-ExecutionMonth] <int> [-ExecutionDayOfMonth]
<int> [-StartTime] <string> [<CommonParameters>]
``` |
| Parameters | • **ExecutionMonth <int>** – Month in which the job runs (1-12).
• **ExecutionDayOfMonth <int>** – Day of month on which the job runs (1-31).
• **StartTime <String>** – Job start time.
• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters". |

Examples	`NewYearlySchedule -ExecutionMonth 1 -ExecutionDayOfMonth 1 -`
	`StartTime "2009-03-31 12:00:00AM"`
	
	This command returns weekly job schdule that runs on the January 1 each year at midnight

Related Links	None

Remove-AllMount

Description	Removes all previously mounted paths. This operation should be performed as a cleanup step after all needed content is retrieved from the mounted path.
Syntax	`Remove-Allmount [<CommonParameters>]`
Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Remove-AllMount` This command removes all mounted paths
Related Links	None

Remove-BackupGroupEntity

Description	The Remove-BackupGroupEntity cmdlet deletes a Backup Group from vRanger. Remove-BackupGroupMember cmdlet is used to remove Members from a Backup Group created using this cmdlet or UI.
Syntax	`Remove-BackupGroupEntity [-BackupGroupId] <GUID>` `[<CommonParameters>]`
Parameters	• **BackupGroupId** <GUID> – The BackupGroup Id • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`$bg = Get-BackupGroupEntity | where {$_.Name -eq 'BG1'}` `Remove-BackupGroupEntity -BackupGroupId $bg.Id`
Related Links	N/A

Remove-BackupGroupMember

Description	The Remove-BackupGroupMember cmdlet is used to remove Members from a Backup Group.
Syntax	`Remove-BackupGroupMember [-BackupGroupId] <GUID>` `[-MemberEntity] <Entity> [<CommonParameters>]`
Parameters	• **BackupGroupId** <GUID> – The BackupGroup Id • **MemberEntity** <Entity> – Inventory entity to remove from Backup Group • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$vm = Get-InventoryEntities | where {$_.Name -eq 'myVM'}`
	
	`Remove-BackupGroupMember -BackupGroupId 5a62765e-6a04-4df4-af54 -`
	`46e5a38ad876 -MemberEntity $vm`

Related Links	N/A

Remove-Catalog

Description	The Remove-Catalog cmdlet is used to delete the catalog data from the database for a list of savepoints.
Syntax	`Remove-Catalog [-Savepoints] <GUID[]> [<CommonParameters>]`
Parameters	• **Savepoints <GUID[]>** – List of Savepoint Ids • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Remove-Catalog -Savepoints 5a62765e-6a04-4df4-af54-46e5a38ad876` Removes the catalog data associated with the savepoint from the database.
Related Links	Get-CatalogSearchData Get-CatalogStatus Set-Cataloging

Remove-DdbStorageUnit

Description	Removes the DDB repository from vRanger and deletes all savepoints in the repository from the disk. The storage unit also is removed from the DD device.
Syntax	`Remove-DdbStorageUnit [-Repository] <Repository> [-Password]` `<String> [<CommonParameters>]`
Parameters	• **Repository** <Repository> – The DDB repository entity to remove. • **Password** <String> – Password for DD Boost user. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Remove-DdbStorageUnit -Repository $repo -Password Password123` This command removes the repository from vRanger and deletes all associated savepoints from the disk. The storage unit also is removed from the DD device.
Related Links	Remove-Repository

Remove-JobTemplate

Description	The Remove-JobTemplate cmdlet removes a JobTemplate from vRanger.
Syntax	`Remove-JobTemplate [-Id] <string> [<CommonParameters>]`
Parameters	• **ID <String>** – The job template ID in GUID. • **<ComMonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Remove-JobTemplate -Id $guid` This command removes the job template with the specified ID from vRanger.
Related Links	Add-BackupGroupEntity Add-RestoreJobTemplate Get-JobTemplate

Remove-LinuxVolume

Description	The Remove-LinuxVolume command is now obsolete and has been replaced with Dismount-LinuxVolume.
Related Links	Dismount-LinuxVolume

Remove-Repository

Description	The Remove-Repository cmdlet removes a repository from vRanger. Optionally, all savepoints in the repository can be deleted from the disk.
Syntax	`Remove-Repository [-Id] <string> [[-RemoveSavePoints] [<Boolean>]] [<CommonParameters>]`
Parameters	• **ID <String>** – The repository ID in GUID. • **RemoveSavePoints [<Boolean>]** – Whether to remove all savepoints in the repositories. 0 or 1. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	------------------EXAMPLE 1-------------------
	
	`Remove-Repository -Id $guid`
	
	This command removes the repository with the specified ID from vRanger. All savepoints associated with the repository remain un-touched.
	
	------------------EXAMPLE 2-------------------
	
	`Remove-Repository -Id $guid -RemoveSavePoints 1`
	
	This command removes the repository withe the specified ID from vRanger. All savepoint files in the repository are removed from disk as well.

Related Links	Add-CIFSRepository Add-SFTPRepository Add-FTPRepository Add-HypervCluster Add-NVSDRepository Get-PhysicalMachineDiskMap

Remove-SavePoint

Description	The Remove-SavePoint cmdlet removes savepoints. If any of the specified savepoint has descendant savepoints, those savepoints are also removed.
Syntax	`Remove-SavePoint [-SavePointsToRemove] <SavePointList> [<CommonParameters>]`
Parameters	• **SavePointsToRemove** <SavePointList> – The list of savepoints to remove. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Remove-SavePoint -SavePointsToRemove $splist` This command removes the list of savepoints specified
Related Links	Get-RepositorySavePoint Remove-Repository

Remove-VirtualAppliance

Description	Removes the virtual appliance configuration from vRanger.
Syntax	`Remove-VirtualAppliance [-VAId] <string> [<CommonParameters>]`
Parameters	• **VAId** <String> – The virtual appliance ID for virtual appliance. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$va = Get-VirtualApplianceConfig -VAName KRN-Vzva-1.1.1`
	`Remove-VirtualAppliance -VAId $va.VirtualApplianceID`
	This command removed a virtual appliance configuration.

Related Links	New-VirtualAppliance Update-VirtualAppliance Get-VirtualApplianceConfig

Remove-VirtualApplianceConfiguration

Description	Removes the virtual appliance configuration from vRanger with an option to delete the VA from host
Syntax	`Remove-VirtualApplianceConfiguration [-VAId] <String> [[-DeleteVM] <Boolean>] [<CommonParameters>]`
Parameters	• **VAId** <String> – VA entity • **DeleteVM** [<Boolean>] – Delete VA from host • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$va = Get-VirtualApplianceConfig -VAName KRN-Vzva-1.1.1`
	`Remove-VirtualAppliance -VAId $va.VirtualApplianceID -DeleteVM 1`
	This command removes a virtual appliance configuration from vRanger and also deletes the VM from Host

Related Links	New-VirtualAppliance Update-VirtualAppliance Get-VirtualApplianceConfig

Run-JobsNow

Description	The Run-JobsNow cmdlet runs jobs specified by their template IDs. Use the **Get-JobTemplate** cmdlet to view the template IDs of configured job templates.
Syntax	`Run-JobsNow [-TemplateIds] <string> [<CommonParameters>]`
Parameters	• **TemplateIds <String>** – The job template IDs. Separate multiple IDs by ','. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Run-JobsNow d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3` This command runs jobs specified by their templates
Related Links	Add-BackupGroupEntity Add-RestoreJobTemplate Get-Job Get-JobTemplate

Run-ReplicationFailover

Description	The Run-ReplicationFailover cmdlet fails over the replication pass to the replica.
Syntax	`Run-ReplicationFailover [-TemplateId] <string> [[-Sync] [<string>]]` `[<CommonParameters>]`
Parameters	• **TemplateId** <String> – Replication job template ID. • **Sync** [<Boolean>] – Whether to sync replica before failing over. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Run-ReplicationFailover -TemplateID d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3` This command runs replication failover based on the job template id.
Related Links	Run-TestReplicationFailover Run-ResumeReplicationFailover

Run-ResumeReplicationFailover

Description	The Run-ResumeReplicationFailover cmdlet resumes test replication failover based on the test replication task id.
Syntax	`Run-ResumeReplicationFailover [-TaskId] <string>` `[<CommonParameters>]`
Parameters	• **TaskID** <String> – Replication test failover task ID. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Run-ResumeReplicationFailover -TaskID d7ebc5c4-7916-44c4-a0c2-` `2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3`
Related Links	Run-TestReplicationFailover Run-ReplicationFailover

Run-TestReplicationFailover

Description	The Run-TestReplicationFailover cmdlet runs test replication failover based on the replication job's template version id.
Syntax	Run-TestReplicationFailover [-TemplateVersionId] <GUID> [[-EnableNetwork] <Boolean>] [<CommonParameters>]
Parameters	• **TemplateVersionID** <String> – Replication job template version ID. • **EnableNetwork** [<Boolean>] – Let replica network be enabled during a test failover. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`Run-TestReplicationFailover -TemplateVersionId d7ebc5c4-7916-44c4-`
	`a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3`
	
	This command runs test failover on a replication job

Related Links	Run-ResumeReplicationFailover Run-ReplicationFailover

Set-Cataloging

Description	The Set-Cataloging cmdlet enables/disables global cataloging flag.
Syntax	`Set-Cataloging [-Catalog] <Boolean> [<CommonParameters>]`
Parameters	• **Catalog** <Boolean> – Boolean value to enable or disable global cataloging • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Set-Cataloging -Catalog 1`
Related Links	Remove-Catalog Get-CatalogSearchData Get-CatalogSearchData

Set-CBTonVM

Description	Enables/Disables block tracking for virtual machines if it is supported
Syntax	`Set-CBTonVM [-VM] <VMEntity> [-CBT] [<Boolean>]` `[<CommonParameters>]`
Parameters	• **VM** <VMEntity> – Virtual machine for which to enable block tracking • **CBT** [<Boolean>] – Enable/Disable Block Tracking • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$vm = Get-InventoryEntities -type virtualmachine | where{$_.Name -`
	`eq 'xyz'}`
	`Set-CBTonVM -VM $vm -CBT 1`
	
	This command enables block tracking on the virtual machine xyz if it is supported on the ESX host.

Related Links	None

Set-LinuxVolume

Description	The Set-LinuxVolume command is now obsolete and has been replaced with Mount-LinuxVolume.
Related Links	Mount-LinuxVolume

Set-MountPath

Description	Sets mount path for a volume. After this operation, the specified volume is mounted to the path specified.
Syntax	`Set-MountPath [-Volume] <Volume> [-Path] <String>` `[<CommonParameters>]`
Parameters	• **Volume <Volume>** – The volume to be mounted • **Path <String>** – The path to which the volume should be mounted to. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Set-MountPath -volume $v -path 'E:\'` This command mount the specified volume in the savepoint to the E:\ drive
Related Links	Get-PhysicalMachineDiskMap Get-RepositorySavePoint Get-RestoreDiskMap Get-VMDKVolume Remove-AllMount

Set-Resources

Description	The Set-Resources cmdlet sets value for vRanger resources.

Syntax	`Set-Resources [[-NumberTasksVRanger] <Int>]`
	`[[-NumberTasksLUN] <Int>] [[-NumberTasksHost] <Int>] [[-`
	`NumberTasksRepository] <Int>] [[-NumberTasksLocal] <Int>] [[-`
	`Timeout] <Int>] [[-MinimumSpaceHost] <Int>] [[-NumberTasksVA]`
	`<Int>] [<CommonParameters>]`

Parameters	• **NumberTasksVRanger** [<Int>] – Maximum number of tasks running on vRanger (>=1) • **NumberTasksLUN** [<Int>] – Maximum number of tasks running off a LUN (1-5) • **NumberTasksHost** [<Int>] – Maximum number of tasks running on a host (1-4) • **NumberTasksRepository** [<Int>] – Maximum number of tasks running per Repository (>=1) • **NumberTasksLocal** [<Int>] – Maximum number of tasks running locally (1-20) • **Timeout** [<Int>] – Task Timeout in hours (1-999) • **MinimumSpaceHost** [<Int>] – Minimum space needed on host (GBs) • **NumberTasksVA** [<Int>] – Maximum number of tasks running on a VA (1-8) • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Set-Resources -NumberTasksVRanger 10 -Timeout 2` Sets the maximum tasks run on vRanger and the Timeout for each task.
Related Links	Get-ConfigOption

Stop-vRangerJob

Description	Cancels a running job.
Syntax	`Stop-vRangerJob [-JobID] <string> [<CommonParameters>]`
Parameters	• **JobID** <String> – Job ID for which to cancel job. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	-------------------**EXAMPLE 1**--------------------
	
	`Stop-vRangerJob -JobId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'`
	This command cancels the job if it is running and returns the success/failure message.
	
	------------------**EXAMPLE 2**--------------------
	
	`$jobs = Get-Job | where {$_.JobState -eq 'Running'}`
	`Stop-vRangerJob $jobs[0].Id`
	This command cancels the first running job in the list and returns the success/failure message.

Related Links	Remove-VirtualApplianceConfiguration Get-Job

Update-BackupJobTemplate

Description	Updates a backup job template. The user also have the option of running the job immediately after the template is updated.
Syntax	`Update-BackupJobTemplate [-JobTemplate] <BackupJobTemplate> [[-JobName] <String>]` `[[-JobDescription] <String>] [[-ExcludedVMList] <String>] [[-IncludedDiskList] <String>]` `[[-TargetRepository] <Repository>] [[-JobSchedule] <Schedule>]` `[[-Flags] <NoFlags	BackupPoweredOnOnly	CheckDestinationFreeSpace	UseCompression	UpdateNotes	EnableGuestQuiescing	EnableABM	UseFiber	PerformNetworkBackupOnFailure	UseCataloging	AllFlags>]` `[[-NumberOfSavePoints] <Int>]` `[[-SpaceSavingTechnologyTypeFlag] <None	Incremental	Differential>] [[-SpaceSavingCountThreshold] <Int>]` `[[-SpaceSavingPercentSizeThreshold] <Int>] [[-NotificationList] <String>] [[-RunJobNow]` `<Boolean>] [[-TransportConfiguration] <TransportConfiguration>]` `[<CommonParameters>]`
Parameters	• **JobTemplate** <BackupJobTemplate> – Existing job template • **JobName** [<String>] – Job name • **JobDescription** [<String>] – Job description • **ExcludedVMList** [<String>] – List of virtual machines to exclude from this job • **IncludedDiskList** [<String>] – List of virtual disks to include in this job • **TargetRepository** [<Repository>] – Repository where the backups are to be stored • **JobSchedule** [<Schedule>] – Schedule on which the job runs • **Flags** [<Nullable\`1>] – Backup flags • **NumberOfSavePoints** [<String>] – Max number of savepoints to store before older ones are removed. • **SpaceSavingTechnologyTypeFlag** [<String>] – Space saving technology to use: differential, incremental, or none. • **SpaceSavingCountThreshold** [<String>] – The number to save before performing a full backup. • **SpaceSavingPercentSizeThreshold** [<String>] – The size to allow as a % of the virtual machine hard disk file before performing a differential backup. • **NotificationList** [<String>] – Email notification list for the job. • **RunJobNow** [<Boolean>] – Whether to run the job now • **TransportConfiguration** [<TransportConfiguration>] – The transport configuration options to use for this job. • **FullBackupSchedule** [<List~1>] – The days to perform a full backup. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters".

Examples	`get-jobtemplate backup |where{$_.jobname -eq 'my first backup job'}|`
	`Update-BackupJobTemplate -JobName 'modified by vapi'`
	
	This command adds a backup job template. Because the RunJobNow flag is not specified, the job is not going to run immediately.

Related Links	Add-BackupJobTemplate Get-Job Get-JobTemplate

Update-GlobalTransportFailover

Description	Updates the default Transport Failover order for replication jobs.
Syntax	`Update-GlobalTransportFailover [-TransportFailover] <string> [<CommonParameters>]`
Parameters	• **TransportFailover** <String> – Transport Failover order separated by ',' to be attempted until the task is successful. Valid types are VA-HotAdd, VA-Network, COS. • **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Update-GlobalTransportFailover -TransportFailover VA-Network, COS` This command updates the default Transport Failover order for replication jobs.
Related Links	Get-GlobalTransportFailover Get-TransportFailover

Update-HypervRestoreJobTemplate

Description	Updates a restore job template. The user also have the option of running the job immediately after the template is updated.
Syntax	`Update-HypervRestoreJobTemplate [-JobTemplate] <RestoreJobTemplate> [-JobName] <String> [[-JobDescription] <String>] [-SavePoint] <SavePoint> [[-JobSchedule] <Schedule>] [[-Flags] <NoFlags	ForceOverwrite	ForcePowerOn	UseFiber	PerformNetworkRestoreOnFailure	AddHypervVirtualMachineToCluster	AllFlags>] [[-NotificationList] <List`1>] [[-RunJobNow] <Boolean>] [[-VMName] <String>] [[-TargetHost] <String>] [[-VmLocation] <String>] [[-SkipDisks] <List`1>] [[-DiskLocations] <List`1>] [[-TargetNetworks] <List`1>] [[-SkipAdapters] <List`1>] [<CommonParameters>]`

Parameters
- **JobTemplate** <RestoreJobTemplate> – Restore job template to update
- **JobName** <String> – Job name
- **JobDescription** [<String>] – Job description
- **SavePoint** <SavePoint> – Savepoint to restore
- **JobSchedule** [<Schedule>] – Schedule on which the job runs. If not set, the job runs on demand.
- **Flags** [<Nullable`1>] – Restore flags. If not set, the job has no flags.
- **NotificationList** [<List`1>] – Email addresses separated by comma.
- **RunJobNow** [<Nullable`1>] – Whether to run the job now. If not set, the job does not run immediately after creation of the template.
- **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used
- **TargetHost** [<String>] – Fully qualified domain name or IP address of the target host.
- **VmLocation** [<String>] – Location on target host where the restored VM is located. If not set, the default value for the host is used.
- **SkipDisks** [<List`1>] – Sequence number (base 0) of the disks to skip from the savepoint. Multiple numbers are separated by commas.
- **DiskLocations** [<List`1>] – Disk locations to use on the target host. Multiple values are separated by comma. If only one is supplied, it is used for all disks on the VM If not set, the default disk location for the host is used for all disks.
- **TargetNetworks** [<List`1>] – Networks to use on the target host. Multiple values are separated by commas. If only one network name is provided, it is used for all adapters on the VM. If not set, the original networks are used when restoring to the original VM.
- **SkipAdapters** [<List`1>] – Sequence number (base 0) of the adapters to skip in the restore job. Multiple numbers are separated by commas.
- **<CommonParameters>** – This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer and OutVariable. For more information, type, "get-help about_commonparameters".

Examples

--------------------EXAMPLE 1--------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$spList = get-repositorysavepoint $repo.id
$mysp = ($spList)[1]
Add-RestoreJobTemplate -Jobname vapirestore  -Savepoint $mysp
```

This command creates a restore job template that restores to the original VM in the savepoint. The job runs on-demand.

--------------------EXAMPLE 2--------------------

```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type HypervHost |
where{$_.name -eq 'dev015.dev.eng.int'}
$targetDiskLocation = 'c:\hyper-virtual hard disks
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
Add-RestoreJobTemplate -jobname vapirestore_totaraget  -savepoint
$mysp -targethost $targethost.name -DiskLocations
$targetDiskLocation -VmLocation 'c:\hyper-v\configuration' -
targetnetworks $targetnw.name
```

This command creates a restore job template that restores to the specified target host. It also specifies disk location, vm location, and the network to use on the target host.

Examples – Continued	-------------------EXAMPLE 3-------------------
| | ```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type HypervHost |
where{$_.name -eq 'dev015.dev.eng.int'}
$targetDiskLocation = 'c:\hyper-virtual hard disks
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$skipdisks = '0'
Add-RestoreJobTemplate -jobname vapirestore_skipdisk  -savepoint
$mysp -targethost $targethost.name -DiskLocations
$targetDiskLocation -VmLocation 'c:\hyper-v\configuration' -
targetnetworks $targetnw.name -SkipDisks $skipdisks
``` |
| | This command creates a restore job template that restores to the specified target host. The first disk (0) in the savepoint is skipped by the restore job. |
| | -------------------EXAMPLE 4------------------- |
| | ```
$repo = get-repository |where{$_.name -eq 'local'}
$sp = get-repositorysavepoint $repo.id
$mysp = ($sp)[0]
$targethost = get-inventoryentities -type EsxHost | where{$_.name -
eq 'dev015.dev.eng.int'}
$targetDiskLocation = 'c:\hyper-virtual hard disks
$targetnw = get-network -host $targethost|where{$_.Name -eq
'Virtual Machine Network'}
$skipnetworks = '0'
Add-RestoreJobTemplate -jobname vapirestore_skipdisk  -savepoint
$mysp -targethost $targethost.name -DiskLocations
$targetDiskLocation -VmLocation 'c:\hyper-v\configuration' -
targetnetworks $targetnw.name -Skipadapters $skipnetworks
``` |
	This command creates a restore job template that restores to the specified target host. The first adapter of the VM in the savepoint is skipped by the restore job.
Related Links	Add-BackupJobTemplate
	Add-RdaRepository
	Update-RestoreJobTemplate
	Get-Job
	Get-JobTemplate

Update-Inventory

Description	Refreshes the inventory.
Syntax	`Update-Inventory [<CommonParameters>]`
Parameters	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`Update-Inventory` This command refreshes the inventory.
Related Links	None

Update-ReplicationJobTemplate

Description	Updates a replication job template. The user also have the option of running the job immediately after the template is updated.
Syntax	`Update-ReplicationJobTemplate [-JobTemplate] <ReplicationJobTemplate> [[-JobName] <String>] [[-JobDescription] <String>] [[-TargetHost] <HostEntity>][[-ReplicateName] <String>] [[-Type] <String>] [[-TargetNetworks] <String>] [[-PrimaryDatastore] <String>] [[-TargetDatastores] <String>] [[-JobSchedule] <Schedule>] [[-TransportFailover] <String>] [[-NotificationList] <String>] [[-Flags] <NoFlags	ReplicatePoweredOnOnly	CheckDestinationFreeSpace	UseCompression	UpdateNotes	EnableGuestQuiescing	EnableABM	AllFlags>] [[-RunJobNow] <Boolean>] [[-IncludedDisk] <String>] [<CommonParameters>]`
Parameters	**JobTemplate** <ReplicationJobTemplate> – Replication job template to update. **JobName** [<String>] – Job name. **JobDescription** [<String>] – Job description. **TargetHost** [<HostEntity>] – Target Host entity to which the VM is replicated **ReplicateName** [<String>] – Replicated VM Name **Type** [<String>] – Replication type. Valid types are Differential and Hybrid **TargetNetworks** <String> – Networks to use on the target host. Multiple values are separated by comma. **PrimaryDatastore** [<String>] – Primary Datastore on target host on which the replicated VM is located. **TargetDatastores** [<String>] – Datastores to use on the target host. Multiple values are separated by comma. **JobSchedule** [<Schedule>] – The recurrence schedule on which the job runs. **TransportFailover** [<String>] – Transport Failover order separated by ',' to be attempted until the tas k is successful. Valid types are VA-HotAdd, VA-Network, COS. **NotificationList** [<String>] – Email notification list for the job. Separate multiple emails by ','. **Flags** [<ReplicationFlags>] – Replication flags: NoFlags, ReplicatePoweredOnOnly, CheckDestinationFreeSpace, UseCompression, UpdateNotes, EnableGuestQuiescing, AllFlags. Multiple flags can be separated by '|' **IncludedDisk** [<int>] – List of disks to include in this job. All VM disks are replicated if none is specified. **RunJobNow** [<Boolean>] – Whether to run the job immediately after the job template is created. **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`get-jobtemplate replication|where{$_.jobname -eq 'my first replication job'}|Update-ReplicationJobTemplate -JobName 'new name'` This command updates a replication job template with a new name.
Related Links	Add-PhysicalMachine Get-Job Get-JobTemplate

Update-RestoreJobTemplate

Description   Updates a restore job template. The user also have the option of running the job immediately after the template is updated.

Syntax
```
Update-RestoreJobTemplate [-JobTemplate] <RestoreJobTemplate> [[-JobName] <String>]
[[-JobDescription] <String>] [[-SavePoint] <SavePoint>] [[-JobSchedule] <Schedule>] [[-Flags] <NoFlags | ForceOverwrite | ForcePowerOn | UseFiber | PerformNetworkRestoreOnFailure | AllFlags>]
[[-NotificationList] <String>] [[-RunJobNow] <Boolean>] [[-VMName] <String>] [[-TargetHost] <String
>] [[-PrimaryDatastore] <String>] [[-SkipDisks] <String>] [[-TargetDatastores] <String>]
[[-TargetNetworks] <String>] [[-SkipAdapters] <String>] [[-TransportConfiguration] <TransportConfiguration>]
[<CommonParameters>]
```

Parameters	• **JobTemplate** <String> – The restore job template to update.

- **JobTemplate** <String> – The restore job template to update.
- **JobName** [<String>] – Job name.
- **JobDescription** [<String>] – Job description.
- **SavePoint** [<SavePoint>] – Savepoint to restore. This parameter is not required if the restore job template contains only one savepoint.
- **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used.
- **JobSchedule** [<Schedule>] – Schedule on which the job runs. If not set, the job runs on demand.
- **Flags** [<RestoreFlags>] – Restore flags. If not set, NoFlag. Restore flags: NoFlags, AllFlags, ForceOverwrite, ForcePowerOn. Multiple flags can be separated by '|'
- **NotificationList** [<String>] – Email notification list for the job. Separate multiple emails by ','.
- **RunJobNow** [<Boolean>] – Whether to run the job now. If not set, the job does not run immediately after creation of the template.
- **VMName** [<String>] – Name of the restored VM. If not set, the VM name in the savepoint is used
- **TargetHost** [<String>] – Fully qualified domain name or IP address of the target host.
- **PrimaryDatastore** [<String>] – Datastore on target host where the restored VM is located. If not set, the original primary datastore is used when restoring to the original VM; otherwise the application randomly selects a datastore attached to the target host.
- **SkipDisks** [<String>] – Sequence number (base 0) of the disks to skip from the savepoint. Disks in savepoint are order by their SCSI controller number and disk number. Multiple numbers are separated by comma.
- **TargetDatastores** [<String>] – Datastores to use on the target host. Multiple values are separated by a comma. If only one is supplied, it is used for all disks on the VM. If not set, the original datastores are used when restoring to the original VM; otherwise the application randomly selects a datastore attached to the target host.
- **TargetNetworks** [<String>] – Networks to use on the target host. Multiple values are separated by a comma. If only one network name is provided, it is used for all adapters on the VM. If not set, the original networks are used when restoring to the original VM; otherwise the application randomly selects a network attached to the target host.
- **SkipAdapters** [<String>] – Sequence number (base 0) of the adapters to skip in the restore job. Multiple numbers are separated by comma.
- **TransportConfiguration** [<TransportConfiguration>] – The transport configuration options to use for this job.
- **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	--------------------EXAMPLE 1--------------------

```
$temp = get-jobtemplate restore |where{$_.jobname -eq 'my restore
job'}
update-restorejobtemplate -jobtemplate $temp -jobname vapirestore
```

This command updates a restore job template name.

--------------------EXAMPLE 2--------------------

```
$savepoint = get-repository |where{$_.name -eq 'local'}|get-
repositorysavepoint |where{$_.VMUUID -eq '12312312312313'}
get-jobtemplate restore |where{$_.jobname -eq 'my restore
job'}|update-restorejobtemplate -savepoint $savepoint
```

This command updates a restore job template to restore a different savepoint.

Related Links	Add-PhysicalMachine
	Get-Job
	Get-JobTemplate

Update-VirtualAppliance

Description	Updates the configuration of a virtual appliance for vRanger.
Syntax	`Update-VirtualAppliance [-VAId] <string> [-IPAddressOverride] <string> [-RootPassword] <string> [<CommonParameters>]`
Parameters	• **VAId** <String>**TargetNetworks [**The virtual appliance ID for virtual appliance to be updated. • **RootPassword** <String>**TargetNetworks [**Root Password for the virtual appliance. • **IPAddressOverride<String> TargetNetworks [**New IP Address for the VA • <CommonParameters>**TargetNetworks [**This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".

Examples	`$va = Get-VirtualApplianceConfig -VAName KRN-Vzva-1.1.1`
	`Update-VirtualAppliance -VAId $va.VirtualApplianceID`
	`-IPAddressOverride 10.6.50.19 -RootPassword vzroot1`
	
	This command updates the password for the virtual appliance.

Related Links	New-VirtualAppliance
	Remove-VirtualAppliance
	Get-VirtualApplianceConfig

Update-VirtualApplianceConfiguration

Description	Updates the virtual appliance configuration from vRanger.

Syntax	`Update-VirtualApplianceConfiguration`
	`[-VirtualApplianceEntity] <VirtualApplianceConfig>`
	`[-IPAddressOverride] <String>] [[-ClusterDefault] <Boolean>][[-`
	`CPUAllocation] <Int>]`
	`[[-MemorySizeInMB] <Int>] [[-ScratchDiskDatastore] <Datastore>] [[-`
	`ScratchDiskSizeInKB] <Int>]`
	`[-NewPassword] <String>][[-UseVAForLFLR] <Boolean>] [[-`
	`DeleteScratchDisk] <Boolean>]`
	`[[-VAReconfigurationStatus] <Boolean>] [<CommonParameters>]`

Parameters	• **VirtualApplianceEntity** <VirtualApplianceConfig> – The VM Entity to reconfigure
	• **IPAddressOverride** [<String>] – IP Address to use for the virtual appliance instead of the one retrieved from the Virtual Center API.
	• **ClusterDefault** [<Boolean>] – If the selected host belongs to a cluster, share the virtual appliance among all hosts in the cluster.
	• **CPUAllocation** [<Int>] – Allocate CPU.
	• **MemorySizeInMB** [<Int>] – Assign memory in MB.
	• **ScratchDiskDatastore** [<Datastore>] – Select a datastore for the ScratchDisk.
	• **ScratchDiskSizeInKB** [<Int>] – Assign ScratchDisk in KB.
	• **NewPassword** [<String>] – New password for virtualAppliance, if this is not specified, the password remains unchanged
	• **UseVAForLFLR** [<Boolean>] – Use virtual appliance for Linux FLR.
	• **DeleteScratchDisk** [<Boolean>] – If set to true, deletes the scratch disk. Default value is false.
	• **VAReconfigurationStatus** [<Boolean>] – Retrieves VA deployment status if set to true .Defaults to false if the parameter is not passed in.
	• **<CommonParameters>** – This cmdlet supports the common parameters: -Verbose, -Debug,-ErrorAction, -ErrorVariable, and -OutVariable. For more information, type, "get-help about_commonparameters".
Examples	`$va = Get-VirtualApplianceConfig -VAName KRN-Vzva-1.1.1 Update-VirtualApplianceConfiguration -VirtualApplianceEntity $va -CPUAllocation 2 -VAReconfigurationStatus 1`
	This command updates the virtual appliance configuration and sets CPUAllocation to 2. Setting VAConfigurationStatus to true retrieves the status of the VA reconfiguration in progress. Default is set to false;
Related Links	New-VirtualAppliance
	Remove-VirtualAppliance
	Get-VirtualApplianceConfig

Using the VI editor

- VI editor overview
- Editing commands
- Saving your changes

This information is a summary of VI functionality, intended to satisfy the basic needs of a user new to VI that just wants to edit the required files for the vRanger virtual appliance configuration. Much more information can be found freely on the internet by searching for "vi editor commands".

> **NOTE:** Linux commands and file names are case sensitive.

VI editor overview

VI is a modal text editor commonly used in Linux and Unix environments. VI operates in either *insert mode* (where keystrokes are entered as part of the document) or *command mode* (where keystrokes are interpreted as commands). The dual modes of VI lets users perform both text editing and command operations without removing their hands from the keyboards home row. The VI editor starts in command mode.

Starting VI

When starting the VI editor, you may create a new file or open an existing file.

- The command **vi** without any file name opens a new file. You need to name and save the file when editing is complete.
- The command **vi** followed by an existing file name opens that file for editing.

Changing modes

Files are opened by the VI editor in command mode, which means that input from the keyboard is treated as VI commands rather than text. Before editing the file, you must change the VI editor to insert mode. The most commands to get into insert mode are listed below. You may also use any of the insertion commands listed in Inserting text.

Table 16. Available commands

Command	Description
a	Enter insert mode; the characters typed are inserted after the current cursor position.
Esc	Switches back to command mode

Editing commands

The sections below contain basic information about navigating the file and editing the text within.

Moving the cursor

Moving the cursor around a file using the VI editor can be a challenge to those not familiar with the commands. Use the commands below as a quick reference.

Table 17. Available commands

Command	Description
h	moves the cursor one space to the left
j	moves the cursor one line down
k	moves the cursor one line up
l	moves the cursor one space to the right
$	moves the cursor to the end of the current line
^	moves the cursor to the beginning of the current line
Enter	moves the cursor to the beginning of the first word on the next line
G	moves the cursor to the end of the file
:*n*	moves the cursor to line *n*
w	moves the cursor to the beginning of the next word
e	moves the cursor to the end of the next word
b	moves the cursor to the beginning of the previous word

Inserting text

Use the commands below to insert text into the open file.

Table 18. Available commands

Command	Description
a	append text starting right of cursor
A	append at the end of the current line
i	insert text starting left of cursor
I	insert text at beginning of the current line
o	open line below cursor, then enter insert mode
O	open line above cursor, then enter insert mode

Deleting text

Use the commands below to delete text from the file.

Table 19. Available commands

Command	Description
x	delete single character
dw	delete word
dd	delete line
s	delete character, remain in insert mode
cw	delete word, remain in insert mode
cc	delete line, remain in insert mode

Table 19. Available commands

Command	Description
u	undo last change
U	undo all changes to current line

Saving your changes

VI does not automatically save your changes. To save your changes, or leave the file without saving changes, use the appropriate command below:

Table 20. Available commands

Command	Description
:w	writes the changes to disk
:wq	writes the changes to disk and quits VI
:q!	quits without saving changes.

About Dell

Dell listens to customers and delivers worldwide innovative technology, business solutions and services they trust and value. For more information, visit www.software.dell.com.

Contacting Dell

Technical support:
Online support

Product questions and sales:
(800) 306-9329

Email:
info@software.dell.com

Technical support resources

Technical support is available to customers who have purchased Dell software with a valid maintenance contract and to customers who have trial versions. To access the Support Portal, go to http://software.dell.com/support/.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. In addition, the portal provides direct access to product support engineers through an online Service Request system.

The site enables you to:

- Create, update, and manage Service Requests (cases)

- View Knowledge Base articles

- Obtain product notifications

- Download software. For trial software, go to Trial Downloads.

- View how-to videos

- Engage in community discussions

- Chat with a support engineer

Free Manuals Download Website

http://myh66.com

http://usermanuals.us

http://www.somanuals.com

http://www.4manuals.cc

http://www.manual-lib.com

http://www.404manual.com

http://www.luxmanual.com

http://aubethermostatmanual.com

Golf course search by state

http://golfingnear.com

Email search by domain

http://emailbydomain.com

Auto manuals search

http://auto.somanuals.com

TV manuals search

http://tv.somanuals.com