

Quest® vRanger® 7.8.5
User's Guide



ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, Foglight, NetVault, vRanger, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. vCenter, vSphere, vMotion, VMware, and ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Active Directory, Hyper-V, SharePoint, SQL Server, Windows, and Windows Server are registered trademarks of Microsoft Corporation in the United States and/or other countries. Data Domain and DD Boost are trademarks or registered trademarks of EMC Corporation in the United States and other countries. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend



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



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



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

Contents

Introduction	10
Quest® vRanger®: at a glance	10
Key benefits	10
Major feature list	11
Target audience	12
Recommended additional reading	13
vRanger overview	14
vRanger overview	14
Installation overview	15
Before installing	15
Downloading an evaluation copy of vRanger	15
Installing vRanger	16
Configuration overview	16
Adding inventory for protection	16
Adding repositories	17
Deploying virtual appliances (VAs)	17
Adding physical servers to inventory	17
Scheduling a backup job	17
Scheduling a replication job	17
Configuring vRanger	19
Configuration overview	19
Configuring vRanger through the Startup Wizard	19
Opening the Startup Wizard	20
Adding and selecting a license	20
Adding vCenter and host information	21
Adding Hyper-V System Center VMM, cluster, and host information	23
Selecting hosts for licenses	24
Adding a physical machine	25
Selecting physical machines for licenses	25
Deploying and configuring the VA from the Startup Wizard	26
Adding a repository	27
Setting up the SMTP server	28
Adding an email address	28
Configuring vRanger manually	29
Adding a vCenter	29
Adding a vCloud Director	30
Adding an ESXi server	30
Adding a Hyper-V System Center VMM, cluster, or host	31
Adding a physical server	33
Adding a repository	34
Managing repository replication	35
Configuring an SMTP server	37

Configuring display language	38
Configuring a virtual appliance (VA)	39
Importing a license file	39
Managing VM host licensing	40
Managing physical machine licensing	40
Managing job configuration settings	41
Enabling cataloging options	42
Managing Transport Failover options	42
Configuring the number of times a backup process tries to run	43
Supplemental instructions: additional repository types	43
Adding an NFS repository	44
Adding a NetVault SmartDisk repository	45
Adding an EMC Data Domain Boost (DD Boost) repository	46
Adding a Quest Rapid Data Access (RDA) repository	48
Understanding the vRanger virtual appliance (VA)	53
The VA scratch disk	53
Deploying the VA by using the Virtual Appliance Deployment Wizard	56
About deploying multiple VAs	59
Changing the VA configuration	59
Configuring VA networking	60
Disabling Distributed Resource Scheduler (DRS) for the vRanger VA	62
Using vRanger	63
vRanger user interface (UI) overview	63
Main toolbar	64
Status Bar	65
My Inventory view	66
My Inventory pane	66
Working Inventory pane	73
My Jobs view	73
My Jobs pane	74
Activity pane	76
My Repositories view	77
My Repositories pane	78
Working Repository pane	79
My Reports view	79
My Reports toolbar	80
My Reports pane	80
Working Reports pane	81
Backup	82
Backup overview	82
Virtual machine (VM) backup	82
Physical machine backup	84
Reviewing backup transport options	84
VMware backup transport options	84
Transport selection order: VM backup	86
Hyper-V backup transport options	87

Physical backup transport options	87
Determining application consistency	88
Application consistency for virtual backups	88
Application consistency for physical backups	90
Understanding retention policies and space-saving technologies	90
Performing optional configurations	92
Enabling VMware Changed Block Tracking (CBT)	92
Adding a custom backup group	93
Configuring VSS tools for VMware VMs	93
Configuring a backup job	94
Starting the Backup Wizard	94
Excluding VMs	96
Including hard disks of physical machines or VMs	96
Selecting a repository	97
Selecting a backup data transport method for VMware machines	97
Configuring backup options	99
Selecting a Retention Policy	100
Setting the recurrence schedule	103
Configuring email notifications	105
Starting the backup job	105
Restore	106
The vRanger restore process	106
Full restore	106
Restoring a file	107
Restoring from a manifest file	107
Performing a full restore for VMware VMs	108
Starting the restore job	114
Performing a full restore for Hyper-V® VMs	114
Performing a full restore for VMware vApps	118
Performing a full restore of a physical machine	123
Starting the restore job	128
Performing an FLR on Windows	128
Performing an FLR using Catalog Search	129
Performing a manual FLR	129
Performing an FLR on Linux	130
Linux FLR limitations	130
Restoring from manifest	131
Performing a full restore from manifest for VMware VMs	131
Performing a full restore from manifest for Hyper-V VMs	137
Performing a full restore from manifest for a physical machine	142
Performing FLR from manifest	148
Replicate VMs	150
Understanding replication of VMs	150
How replication works	151
Replication modes	153

Replication with Changed Block Tracking (CBT)	153
Differential replication	154
Configuring transport failover options	154
Pre-seeding replication jobs	154
Creating a replication job	155
Selecting the source VM	155
Selecting a target host	156
Selecting the replication mode	156
Selecting replication network mapping	156
Selecting hard disks	157
Selecting replication options	157
Setting the recurrence schedule	158
Configuring email notifications	160
Starting the replication job	160
Using replication failover	160
Performing Test Failover with networking disabled	161
Performing Test Failover with networking enabled	161
Performing a failover	162
Performing failback	163
Reports	165
About vRanger reports	165
Generating a report using the Report Wizard	165
Integrating and monitoring vRanger	170
Integrating vRanger: overview	170
Performing vRanger sweep-to-tape using Quest NetVault Backup	170
Using deduplication with Quest NetVault SmartDisk and vRanger	177
Using Data Domain and vRanger repositories	177
Monitoring vRanger	177
About event monitoring configurations	178
Using the vRanger Console	181
vRanger Console overview	181
Configuring the vRanger Console	181
vRanger vAPI Cmdlets	182
Inventory management	182
Configuration	182
Job management	182
Repositories	183
Backup	183
Restore	183
Replication	183
Cmdlet help	183
vAPI Cmdlet details	184
Add-BackupGroupEntity	184
Add-BackupJobTemplate	185

Add-CIFSRepository	186
Add-DdbReplicationRepository	187
Add-DdbRepository	188
Add-EsxHost	189
Add-HypervCluster	189
Add-HypervHost	190
Add-HypervRestoreJobTemplate	190
Add-NFSRepository	193
Add-NVSDRepository	194
Add-PhysicalMachine	194
Add-RdaRepository	195
Add-ReplicationJobTemplate	196
Add-RestoreFromManifestJobTemplate	197
Add-RestoreJobTemplate	200
Add-VirtualAppforLinuxFLR	204
Add-VirtualAppforLinuxFLRVA	204
Add-VirtualCenter	204
Disable-Job	205
Dismount-LinuxVolume	205
Enable-Job	205
Get-AddressBook	206
Get-BackupGroupEntity	206
Get-CatalogSearchData	207
Get-CatalogStatus	207
Get-ConfigOption	208
Get-Connection	208
Get-CurrentTemplateVersionID	209
Get-Datstore	209
Get-GlobalTransportFailover	209
Get-InventoryEntity	210
Get-IsInventoryRefreshing	210
Get-Job	211
Get-JobTemplate	212
Get-MonitorLog	212
Get-Network	213
Get-PhysicalMachineDiskMap	213
Get-Repository	213
Get-RepositoryJob	214
Get-RepositorySavePoint	214
Get-RestoreDiskMap	215
Get-SavepointDisk	215
Get-SavepointManifest	216
Get-Savepoints	216
Get-TransportFailover	216
Get-VirtualApplianceConfig	217

Get-VirtualApplianceDeploymentStatus	217
Get-VirtualApplianceReconfigStatus	218
Get-VirtualMachinesUnderInventory	218
Get-VmDisk	219
Get-VMDKVolume	219
Install-VirtualAppliance	220
Mount-LinuxVolume	222
New-BackupFlag	222
New-BackupGroupMember	223
New-Daily Schedule	224
New-EmailAddress	224
New-IntervalSchedule	225
New-MonthlySchedule	225
New-ReplicationFlag	226
New-RestoreFlag	227
New-SMTPServer	227
New-TransportConfiguration	228
New-VirtualAppliance	229
New-WeeklySchedule	229
New-YearlySchedule	230
Remove-AllMount	230
Remove-BackupGroupEntity	230
Remove-BackupGroupMember	231
Remove-Catalog	231
Remove-DdbStorageUnit	232
Remove-JobTemplate	232
Remove-LinuxVolume	232
Remove-Repository	233
Remove-SavePoint	233
Remove-VirtualAppliance	234
Remove-VirtualApplianceConfiguration	234
Run-JobsNow	234
Run-ReplicationFailover	235
Run-ResumeReplicationFailover	235
Run-TestReplicationFailover	236
Set-Cataloging	236
Set-CBTonVM	236
Set-LinuxVolume	237
Set-MountPath	237
Set-Resources	237
Stop-vRangerJob	238
Update-BackupJobTemplate	239
Update-GlobalTransportFailover	240
Update-HypervRestoreJobTemplate	240
Update-Inventory	242

Update-ReplicationJobTemplate	242
Update-RestoreJobTemplate	244
Update-VirtualAppliance	246
Update-VirtualApplianceConfiguration	246
About us	248
We are more than just a name	248
Our brand, our vision. Together.	248
Contacting Quest	248
Technical support resources	248

Introduction

- [Quest® vRanger®: at a glance](#)
- [Key benefits](#)
- [Major feature list](#)
- [Recommended additional reading](#)
- [Recommended additional reading 7.8.5](#)

Quest® vRanger®: at a glance

vRanger provides a simple, fast, and scalable data- protection solution that deploys seamlessly into virtual environments. vRanger provides high-speed backup and recovery of VMware®, 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 virtual machines (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 supports enterprise deduplication and delivers disk-to-disk backup with Quest® QoreStor™, Quest® DR Series System appliances, EMC® Data Domain®, and Quest® NetVault™ SmartDisk to reduce backup storage costs and improve backup time with client-side deduplication.

vRanger scales with your virtual environment by maximizing resources through distributed processing, while simplifying management with central command and control.

vRanger capabilities include:

- Protect entire VMware and Hyper-V environments in minutes.
- Deliver high-speed, storage-thrifty backup and 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

vRanger benefits include:

- Exploits VMware vSphere® performance and scalability enhancements.
- Provides high-speed, resource-efficient backup, replication, and recovery of VMware® and Hyper-V® VM images.

- Delivers maximum storage savings when paired with Quest® QoreStor™, Quest® DR Series System appliances, EMC® Data Domain®, and Quest® NetVault™ SmartDisk
- Supports EMC Data Domain Boost and Quest QoreStor and DR Series RDA for optimized deduplication and replication.
- Operates as an easy-to-deploy, low-resource consumption virtual appliance (VA) for low-impact scalability.
- Improves VM backup and recovery performance with HotAdd support.
- Optimizes use of critical resources in virtual environments, including hypervisor hosts, networks, and storage.
- Includes a catalog for fast file searching and file-level restore.
- Provides protection for the largest VM deployments.
- Offers agent-less, maintenance-free deployment.
- Delivers proven reliability.
- Provides unmatched performance and scalability.
- Comes with world-class service and support.

Major feature list

vRanger backs up and recovers physical environments with blazing speed, a minimum of backup storage, and can restore Windows servers and even individual files and folders. With vRanger, you get comprehensive protection for both your virtual and physical environments that you can manage from one simple interface.

- **VMware® Ready™ Certified for vSphere®:** Exploits the latest in vSphere 6.7 performance and scalability improvements.
- **VMware vSphere / ESXi / vCloud / Virtual SAN, Virtual Volumes, Storage Profiles, and DRS cluster support** – Provides a wide range of backup and recovery and disaster recovery (DR) capabilities for VMware virtual infrastructures.
- **Microsoft Hyper-V 2019:** Protect the latest Hyper-V versions.
- **Virtual appliance (VA) architecture:** Through centralized and wizard-driven deployment and administration from the vRanger server, delivers scalable and cost-effective distributed data handling and throughput through VAs.
- **Disk-to-disk backup and deduplication:** Supports Quest DR Series system Rapid Data Access (RDA) for optimized client-side deduplication and replication.
- **EMC® Data Domain® (DD Boost™) support:** Supports distributed deduplication with EMC Data Domain appliances using the DD Boost API.
- **Disk-to-disk backup and deduplication:** Offers a Quest NetVault SmartDisk - vRanger Edition add on for enterprise-class deduplication that reduces backup storage footprints by up to 90 percent.
- **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.
- **Change block tracking (CBT):** Eliminates the time required to scan for changed blocks in guest images on vSphere hypervisor systems to improve the speed of backup and replication jobs.
- **Instant file-level restore (FLR) for Windows and Linux®:** Lets you restore a single file from a backup image in the repository through a one-step process.

- **Patented native, full catalog capability:** 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 locates backup repository files to be restored.
- **One-step catalog recovery:** Provides advanced search (including wildcards) and right-click recovery selection directly from the management console to speed restore of VMs, savepoints and hosts with native catalog.
- **Advanced Encryption Standard (AES)-256:** Secures protected images block-by-block on the VMware host as they are read so they are also secure over the network and in the backup repository.
- **Full, incremental, and differential backup:** Enables a complete backup cycle for protected images that is optimized for speed and resource efficiency.
- **Full and incremental replication:** Provides the full range of options required to replicate VMware VMs over LANs and WANs efficiently.
- **Replication:** Offers on and off-site options for flexible, reliable disaster recovery.
- **Supports backup and replication of encrypted VMware vSphere VMs:** Provides the full range of options required to backup, restore, replicate, failover, and failback encrypted virtual machines.
- **Physical Windows Server® support:** Supports backup and recovery of Windows physical servers, files, and folders.
- **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.
- **VMware vSphere® 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.
- **Agent-less job execution, patent-pending:** Uses binary injection at run time on VMware ESXi hosts; eliminates burdens of license tracking and maintenance upgrades.
- **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.
- **Job multi-streaming:** Lets you run multiple backup, restore, and replication jobs simultaneously for dramatic improvement in performance times.
- **Direct-to-target (D2T) architecture:** Distributes job execution and movement of data with optimal efficiency to improve data protection and ensure seamless scalability.
- **FIPS-compatible:** Address government data security requirements with vRanger Crypto Module, compatible with FIPS 140-2 Level.
- **Inline data validation:** Tests integrity of captured data on the source block-by-block as it is read; ensures image recovery from the backup repository and replica-image usability.
- **Remote management:** Lets you manage data protection jobs through a central console over LAN and WAN connections for control of all systems and sites in an environment.
- **PowerShell® access:** Automate scripts for protection jobs, reducing administrative burdens and human error.

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:

- *Quest 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.
- *Quest vRanger Integration Guide for EMC® Data Domain® Boost (DD Boost™)*: This guide describes how to configure vRanger to work with a Data Domain Boost repository to achieve source-side deduplication of backup data.
- *Quest vRanger Integration Guide for Quest DR Series Disk Backup Appliance*: This guide describes how to configure vRanger to work with a Quest DR series appliance to achieve deduplication of backup data.
- *Quest vRanger Integration Guide for Quest 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.
- *Quest vRanger Integration Guide for Quest QoreStor* - This guide describes how to configure vRanger to work with QoreStor to achieve deduplication of backup data.

vRanger overview

- [vRanger overview](#)
- [Installation overview](#)
- [Configuration overview](#)

vRanger overview

vRanger provides backup, recovery, and backup management 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 include:

- **Protect entire VMware® and Hyper-V® environments in minutes:** Deploy vRanger's simple installation wizard and auto-discovery tool to locate both VMware and Hyper-V assets, and safeguard entire virtual environments in minutes. Leverage vRanger's agentless architecture to reduce time spent managing virtual backup, restore and recovery operations, and to detect and back up new VMs automatically for complete protection.
- **Deliver high-speed, storage-thrifty backup and restore for VMware, Hyper-V, and physical Windows® servers:** Employ the fastest I/O path available to back up and restore VMware, Hyper-V, and physical Windows servers quickly, with minimal disruption to production systems. vRanger's patented Active Block Mapping feature reduces data streams, network utilization and backup storage requirements. Leverage client-side and target-side deduplication options with the support for the Quest DR disk-based backup appliance and EMC™ Data Domain integration to dramatically shrink backup storage costs.
- **Quickly locate and restore individual files from virtual and physical backups:** Locate and restore specific versions of VMs, files, and folders in minutes. Use vRanger's integrated catalog and wildcard search features to find individual files quickly, whether they are located within VMware, Hyper-V, or physical Windows Server backups. Manage virtual and physical servers from a single, easy-to-use user interface (UI), minimizing the effort, complexity and cost of backing up your data.
- **Replicate key VMs for offsite disaster recovery preparedness:** Replicate VMware virtual machines both on-site and at remote locations for flexible and efficient disaster recovery preparedness. Combine backup and replication jobs to ensure that you meet your organization's recovery time and recovery point objectives. Perform fast VM failover and recovery at all your sites, no matter where they are located.
- **Deploy scalable data protection for even the largest virtual infrastructures:** Leverage parallel job execution capabilities to complete backup and recovery operations faster, and use resource governing to ensure optimal production application performance. Deploy the vRanger virtual appliance (VA) to distribute VMware workloads for better scalability. Use vRanger's LAN-free mode to move backup and restore operations off local-area networks (LANs) for improved virtualized application performance.
- **Use Advanced Encryption Standard (AES)-256:** vRanger uses AES-256 encryption to encrypt backup archives and replication traffic. In vRanger, AES is a symmetric key algorithm that uses a 256-bit encryption key. This combination is approved for the encryption of secure government data, and complies with the Federal Information Processing Standard (FIPS).

vRanger requires credentials for several connection types; namely local machine credentials, vCenter and host credentials, and credentials for the vRanger database. These credentials are encrypted using AES-256 encryption before they are written to the vRanger database or Windows registry.

Encryption for vRanger backups is configured at the repository level, meaning that rather than selecting encryption for each job or task, you configure that job to write to an encrypted repository. For each backup task, vRanger reads the encryption status of the target repository. If encryption is selected, the data blocks are read and encrypted before they are sent to the repository.

During replication tasks, vRanger sends data over a secure shell (SSH) tunnel encrypted with AES-256 encryption.

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 VM meeting the requirements specified in the *Quest vRanger Installation/Upgrade Guide*.

- [Before installing](#)
- [Downloading an evaluation copy of vRanger](#)
- [Installing vRanger](#)

Before installing

vRanger can be installed in various 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 (VA) deployment

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

For more information, see the *Quest vRanger Installation/Upgrade Guide*.

Downloading an evaluation copy of vRanger

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

To download an evaluation copy of vRanger:

- 1 From an internet-connected PC, access <http://quest.com/products/vranger/>.
- 2 Click **Download Free Trial**.
- 3 Register or log in as appropriate.
- 4 Accept the **License Agreement**, and click **Submit**.
- 5 Select **vRanger**, and save the file to the desired location.

Installing vRanger

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

To install vRanger:

- 1 Start 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 now.

The installation completes.

For more information, see the *Quest vRanger 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 starts the first time the application is opened. For more information, see [Configuring vRanger](#).

The following topics describe the primary configurations you need to make.

- [Adding inventory for protection](#)
- [Adding repositories](#)
- [Deploying virtual appliances \(VAs\)](#)
- [Adding physical servers to inventory](#)
- [Scheduling a backup job](#)
- [Scheduling a replication job](#)

Adding inventory for protection

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

- One or more VMware® vCenter Servers™.
- One or more Hyper-V® clusters.
- One or more VMware® vCloud Director® servers.
- One or more physical servers.
- Individual Hyper-V® or VMware® 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 can exclude hosts from the vRanger inventory, which also excludes them from licensing.

Adding repositories

Repositories are where vRanger stores the savepoints created by each backup job. You can create a repository from a standard Network File System (NFS) or Common Internet File System (CIFS). You may also deduplicate vRanger backups by sending them to a Quest DR RDA, EMC® Data Domain® (DD Boost™), or Quest NetVault SmartDisk repository. To add a repository, you need the name of the server and share, and an account with access to that share. If you are using Quest RDA or DD Boost repositories, you can also set up repository replication with another repository of the same type. For more information on replication, see [Managing repository replication](#).

Deploying virtual appliances (VAs)

vRanger uses a VA for replication to and from VMware® ESXi™ servers, for Linux® FLR, and optionally for backups and restores. Before fully utilizing vRanger, you need to deploy and configure VAs to your ESXi servers. For more information, see the *Quest vRanger Installation/Upgrade Guide*.

Adding physical servers to inventory

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

Scheduling a backup job

With a VMware® vCenter™ and a Repository configured, the next step is to schedule a backup job.

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.
- 8 Configure Retention Policies.
- 9 Schedule the backup job or set to run On Demand.

For more information about the Backup Wizard or available options, see [Backup](#).

Scheduling a replication job

With a VMware® vCenter™ and a VA or VAs—if necessary—configured, the next step is to configure a replication job.

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 <VMname>**.
- 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 or disks 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 VMs](#).

Configuring vRanger

- [Configuration overview](#)
- [Configuring vRanger through the Startup Wizard](#)
- [Configuring vRanger manually](#)
- [Supplemental instructions: additional repository types](#)
- [Understanding the vRanger virtual appliance \(VA\)](#)

Configuration overview

Before configure vRanger, confirm that you have access to all user names, passwords, and root passwords for the VMware® ESXi™ 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 starts 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 box to manage your settings. Although the Configuration Options dialog box 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 box 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 set up an email Simple Mail Transfer Protocol (SMTP) server 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 through the Startup Wizard, complete the following tasks and processes:

- [Opening the Startup Wizard](#)
- [Adding and selecting a license](#)
- [Adding vCenter and host information](#)
- [Adding Hyper-V System Center VMM, cluster, and host information](#)
- [Selecting hosts for licenses](#)
- [Adding a physical machine](#)
- [Selecting physical machines for licenses](#)
- [Deploying and configuring the VA from the Startup Wizard](#)

- [Adding a repository](#)
- [Setting up the SMTP server](#)

Opening the Startup Wizard

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

To start the wizard:

- 1 Do one of the following:
 - If you just installed the software, the Startup Wizard automatically starts.
 - If the software was already installed, on the Main toolbar, click **Tools > Startup Wizard**.
- 2 In the **Startup Wizard** dialog box, 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 and a trial license.

To add and select a license:

- 1 Click **Add New License**.
- 2 Select the license, and click **OK** to import it.

CAUTION: Adding a license with fewer features or CPUs than the current license could result in the loss of functionality, or in some hosts or machines becoming unlicensed.

- 3 Review the information in the **License information for <licensefile>** section to make sure that it displays the expected allotments.

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

- 4 Click **Save**.

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 **Licenses Allowed** and **Licenses Used** counts.

- 5 Click **Next**.

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 On the Main toolbar, click **Tools > Options**.
- 2 In 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 applied.
- 6 For the **Virtual Machine Backup**, **Physical Machines Backup**, and **Virtual Machine Replication** nodes, confirm your license allocation.

Adding vCenter and host information

The VirtualCenter and Host Information page appears. Before vRanger can add source objects to the My Inventory pane, you must configure a connection to a VMware® vCenter™ or a VMware® ESXi™ Server, or both.

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

- **Adding a vCenter:** Complete this task to add a vCenter and all associated hosts.
- **Entering host credentials:** Complete this task to add credentials to hosts that appeared after adding a vCenter.
- **Adding hosts:** Complete this task to add individual hosts not associated with a vCenter.

Adding a vCenter

Complete the following steps to add a VMware® 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** field, enter the FQDN or IP address of the vCenter Server™.
 - 3 In the **User Name** field, enter the name of an account with privileges on the vCenter Server.
For the required permissions for a vRanger vCenter account, see the *Quest vRanger Installation/Upgrade Guide*.
- i | NOTE:** The user name for the vCenter credential should be entered in the **username@domain** format, rather than **domain\username**. Sometimes, the domain might not be required. Avoid special characters in the user name.



If these credentials are changed in the future, you must restart the Quest vRanger Service to recognize the changes.
- 4 In the **User Password** field, enter the password for the preceding account.

i | NOTE: Avoid special characters in the password.
 - 5 In the **Port Number** field, enter the port to use for communication—the default is 443.
 - 6 Click **Connect**.

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

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: .

CAUTION: vCenter credentials are sufficient for operations that use only the vStorage API.

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.

TIP: To select all hosts, select the upper-most check box.

- 2 Click **Edit**.

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

- 3 In the **User Name** field, 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** field, enter the password for the preceding account.

- 5 In the **Root Password** field, enter the root password for the selected hosts.

If your hosts use different root passwords, configure them in like groups based on the root password.

- 6 In the **Port Number** field, enter the port to use for communication—the default 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 alarm appears, 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: VMware® 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 VMware® vCenter™, you can add them individually.

To add hosts:

- 1 In the **Hosts** section, click **Add**.
- 2 In the **DNS Name or IP** field, enter the FQDN or IP address of the Host.
- 3 In the **User Name** field, 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** field, enter the password for the preceding account.
- 5 In the **Root Password** field, enter the password for the root account.
- 6 In the **Port Number** field, enter the port to use for communication—the default is 22.
- 7 Click **Connect**.
The host appears 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 appears. Before vRanger can add source objects to the My Inventory pane, you must configure a connection to a Hyper-V® System Center Virtual Machine Manager (VMM), cluster, or host.

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

- [Adding a Hyper-V System Center VMM](#)
- [Adding a Hyper-V cluster](#)
- [Adding hosts](#)

Adding a Hyper-V System Center VMM

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

- i** | **IMPORTANT:** Before adding the System Center VMM connection, you must install the System Center VMM Administrator Console on the vRanger machine.
- IMPORTANT:** Only one version of the SCVMM Console may be installed on the vRanger machine. vRanger does not support simultaneous operations against two different versions of SCVMM from the same 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** field, enter the FQDN or IP address of the Hyper-V System Center VMM.
- 3 In the **User Name** field, enter the name of an account with domain administrator privileges on the System Center VMM.
- 4 In the **User Password** field, enter the password for the preceding account.
- 5 If this VMM 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** field, 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 is 8081.
 - In the **SCVMM Port Number** field, enter the port you want vRanger to use to communicate with the System Center VMM server—the default is 8100. Click **Connect**.

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

- 7 Click **Next**.

The Hyper-V System Center VMM appears in the **Hyper-V System Center Virtual Machine Managers** section. The hosts managed by that VMM appear 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** field, enter the FQDN or IP address of the Hyper-V cluster.
- 3 In the **User Name** field, enter the name of an account with domain administrator privileges on the cluster.
- 4 In the **User Password** field, enter the password for the preceding account.
- 5 If this cluster is new, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.
- 6 In the **Port Number** field, 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 is 8081.
- 7 Click **Connect**.
- 8 Click **Next**.

The Hyper-V cluster appears in the **Hyper-V Clusters** section. The hosts managed by that cluster appear 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 **Hosts** section, click **Add**.
- 2 In the **DNS Name or IP** field, enter the FQDN or IP address of the Host.
- 3 In the **User Name** field, enter an account for the host.
- 4 In the **User Password** field, enter the password for the preceding account.
- 5 If this host is new, or if you have removed the vRanger Hyper-V® agent, select **Install agent on host**.
- 6 In the **Port Number** field, 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 is 8081.
- 7 Click **Connect**.

The host appears in the **Hosts** section.

Selecting hosts for licenses

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

i | **NOTE:** VM licenses include all VMs. 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 **Backup or Replication** column, select or clear hosts until you have used all the available licenses, or until all your hosts are protected.
- 3 Click **Save**.
- 4 Click **Next**.

Adding a physical machine

The Physical Machine Information page appears. 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 In the **DNS Name or IP** field, enter the FQDN or IP address of the server.
- 3 In the **User Name** field, enter an account for the server.
- 4 In the **User Password** field, enter the password for the preceding account.
- 5 If this server is new, or if you have removed the vRanger agent, select **Install agent on machine**.
In the **Agent Location** field, enter the preferred directory—on the physical machine—to which the physical client should be installed. The default is **C:\Program Files\Dell\VRangerPhysicalClient**.
- 6 In the **Port Number** field, 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 is 51000.
- 7 Click **Connect**.

The server appears in the My Inventory pane. You may also create a backup group to combine multiple physical servers into one backup job. For more information, see [Adding a custom backup group](#).

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. The **Licensed Physical Machine** section lists the physical machines in the vRanger inventory, and indicates whether each machine is licensed for backup.
- 2 In the **Physical Backup** column, select or clear machines until you have used all the available licenses, or until all your machines are protected.

- 3 Click **OK**.

Deploying and configuring the VA from the Startup Wizard

vRanger uses a virtual appliance (VA) for Linux® file-level restore (FLR), standard backup and restore, 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 box 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 VA from the Startup Wizard:

- 1 On the **Virtual Appliance Information** page of the Startup Wizard, start 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 VA by using the Virtual Appliance Deployment Wizard](#).
- 3 To add a VA configuration, click **Add**, and then complete the following steps:
 - a In the **Add Virtual Appliance Configuration** dialog box, select a VA from the inventory tree.
 - b Under **Virtual Appliance Properties**, do any of the following:
 - Select **Override IP Address**, and then enter a new IP address in the **IP Address** field.
 - Enter a root password for the VA in the **Root Password** field.
 - 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 VA, select a VA from the list, and then click **Edit**.

In the **Modify Virtual Appliance Configuration** dialog box, 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 VA configuration](#).

- 5 To delete a VA, do the following:
 - a Select a VA from the list.
 - b Click **Remove**.
The Removing VA dialog box appears.
 - c Select the job you want to remove, and click **OK**.
 - d In the **Confirm Delete** dialog box, click **OK**.

If you want to 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 FLR, select the VA you want to use from the drop-down list.
- 7 Click **Next**.

Adding a repository

vRanger uses repositories to store backup archives. vRanger supports the following options for repository connections.

i | NOTE: Any storage device or appliance that supports CIFS or NFSv3 is expected to be compatible with vRanger.

Windows Share (CIFS)

- **NFS** (version 3)
- **NetVault SmartDisk (NVSD)**: The Quest 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://quest.com/products/netvault-smartdisk/> or the *Quest vRanger Integration Guide for Quest NetVault SmartDisk*.
- **EMC Data Domain Boost (DDB)**: Integrating EMC Data Domain Boost (DD 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 DD Boost, see <http://www.emc.com/data-protection/data-domain/data-domain-boost.htm> or the *Quest vRanger Integration Guide for EMC Data Domain Boost (DD Boost)*.
- **Quest Rapid Data Access (RDA)**: Provided by the Quest DR Series appliances—purpose-built, disk backup appliances that use Quest deduplication technology to improve backup and recovery processes. For more information on Quest DR Series appliances, see <http://quest.com/products/dr-series-disk-backup-appliances/> or the *Quest vRanger Integration Guide for Quest DR Series Disk Backup Appliance*.

i | NOTE: Quest RDA and EMC DD Boost repositories support repository replication. For more information, see [Managing repository replication](#).

The following procedure 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** field.
This value appears in the **My Repositories** pane.
- 3 Populate the **Description** field.
- 4 Enter a **username** and **password** in the appropriate fields.
- 5 Select a **Security Protocol** from the drop-down list: NTLM (Default), or NTLM v2
- 6 In the **Server** field, type the UNC path to the preferred repository directory.
Alternatively, enter a partial path, and click **Browse** to find the target directory.

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

! | CAUTION: If you use the **Encrypt all backups to this repository** feature, make certain to retain the password you enter. There is no back-door or administrator-level password. If the password is unknown, the backups are not usable.

- 7 If you want these backups password-protected, select **Encrypt all backups to this repository**, enter a Password for the encrypted repository, and then confirm the password by re-entering it.

i | NOTE: Encryption is *not* supported for NetVault SmartDisk and DD Boost repositories.

- 8 Enter a Password for the encrypted repository, and 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 box.

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 10 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
 - **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of 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** field, enter the DNS name or IP of an SMTP server.
 - 2 In the **Port** field, enter the port to use for communication—the default is 25.
 - 3 If your SMTP requires use of authentication or encryption, or both, click **Advanced**, and complete the following in the **SMTP Server Advanced** dialog box:
 - a To use authentication, select the **SMTP server requires authentication** check box, and enter the applicable user name and password. If applicable, include the domain followed by a backward slash before the name.
 - b To use SSL, select the **Connect using SSL** check box.
 - c Click **OK**.
 - 4 In the **From Address** field, enter the email address that vRanger should use.
- i | NOTE:** Make sure that you use a valid email account registered to the SMTP server that you added.
- 5 Click **Send Test Email**.
 - 6 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** fields.
- 3 Click **OK** to add the information to the Email Addresses list.
- 4 Repeat [Step 1](#) through [Step 3](#) to include additional contacts.
- 5 Click **Finish**.

Configuring vRanger manually

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® VMs, 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 vCenter](#)
- [Adding a vCloud Director](#)
- [Adding an ESXi 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 \(VA\)](#)
- [Importing a license file](#)
- [Managing VM host licensing](#)
- [Managing physical machine licensing](#)
- [Managing job configuration settings](#)
- [Enabling cataloging options](#)
- [Managing Transport Failover options](#)
- [Configuring the number of times a backup process tries to run](#)

Adding a vCenter

Complete the steps in the following procedure to add a VMware® vCenter™ for vRanger backup.

To add a VirtualCenter:

- 1 In the **My Inventory** view, click the **VMware** tab.
- 2 Click the **Add** icon, and then click **Virtual Center**.
The **Add VirtualCenter Credentials** dialog box appears.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the VMware® vCenter Server™.
- 4 In the **User Name** field, enter the name of an account with privileges on the vCenter Server.

For the required permissions for a vRanger vCenter account, see the topic on configuring vCenter permissions in the *Quest vRanger Installation/Upgrade Guide*.

NOTE: The user name for the vCenter credential should be entered in the **username@domain** format, rather than **domain\username**. Sometimes, the domain name might not be required. Avoid special characters in the user name.

If these credentials are changed in the future, you must restart the vRanger Service to recognize the changes.

- 5 In the **User Password** field, enter the password for the preceding account.

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 appears in the **My Inventory** pane.

Adding a vCloud Director

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

To add a vCloud Director:

- 1 In the **My Inventory** view, click the **vCloud** tab.
- 2 Click the **Add** icon, and then click **vCloud Director**.
The **Add vCloud Director Credentials** dialog box appears.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the vCloud Director server.
- 4 In the **User Name** field, enter the name of an account with administrator privileges on the vCloud Director server.
- 5 In the **User Password** field, enter the password for the preceding account.

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 appears in the **My Inventory** pane.

You are prompted to add credentials for each VMware® vCenter Server™ attached to the vCloud Director.

- 7 In the **Add vCloud Director Credentials** dialog box, 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, and click **Connect** to authenticate.
- 9 Click **Done**.

Adding an ESXi server

Complete the steps in the following procedure to add a VMware® ESXi™ server to vRanger protection.

To add an ESXi server:

- 1 In the **My Inventory** view, click the **VMware** tab.
- 2 Click the **Add** icon, and then click **ESX Host**.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the Host.
- 4 In the **User Name** field, 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** field, enter the password for the preceding account.
- 6 In the **Root Password** field, enter the password for the root account.
- 7 Click **Connect**.

If the ESXi server is already part of a vCenter added to vRanger, the host is added to the tree structure under the vCenter.

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

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

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

To add a Hyper-V host:

- 1 In the **My Inventory** view, click the **Hyper-V** tab.
- 2 Click the **Add** icon, and then click **Hyper-V Host**.
The **Add Credentials** dialog box appears.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the Host.
- 4 In the **User Name** field, enter an account for the host.
- 5 In the **User Password** field, enter the password for the preceding account.
- 6 If this host is new, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.

In the **Port Number** field, 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 is 8081.


To add a Hyper-V System Center VMM:

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

IMPORTANT: Only one version of the SCVMM Console may be installed on the vRanger machine. vRanger does not support simultaneous operations against two different versions of SCVMM from the same vRanger machine.

- 1 In the **My Inventory** view, click the **Hyper-V** tab.
- 2 Click the **Add** icon, and then click **Hyper-V System Center VMM**.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the Hyper-V System Center VMM.

- 4 In the **User Name** field, enter the name of an account with domain administrator privileges on the System Center VMM.
- 5 In the **User Password** field, enter the password for the preceding account.
- 6 If this VMM is new System Center VMM, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.
- 7 Configure the ports as follows:
 - In the **Agent Port Number** field, 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 is 8081.
 - In the **SCVMM Port Number** field, enter the port you want vRanger to use to communicate with the System Center VMM server—the default is 8100. Click **Connect**.

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

- 8 Click **Next**.

The VMM appears in the **Hyper-V System Center Virtual Machine Managers** section. The hosts managed by that VMM appear 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 **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** field, enter the FQDN or IP address of the Hyper-V cluster.
- 4 In the **User Name** field, enter the name of an account with domain administrator privileges on the cluster.
- 5 In the **User Password** field, enter the password for the preceding account.
- 6 If this cluster is new, or if you have removed the vRanger Hyper-V agent, select **Install agent on host**.
- 7 In the **Port Number** field, 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 is 8081.

- 8 Click **Connect**, and then click **Next**.

The Hyper-V cluster appears in the **Hyper-V Clusters** section. The hosts managed by that cluster appear 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 **Hosts** section, click **Add**.
- 2 In the **DNS Name or IP** field, enter the FQDN or IP address of the Host.
- 3 In the **User Name** field, enter an account for the host.
- 4 In the **User Password** field, enter the password for the preceding account.
- 5 If this host is new, or if you have removed the vRanger Hyper-V® agent, select **Install agent on host**.

- 6 In the **Port Number** field, 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 is 8081.

- 7 Click **Connect**.

The host appears in the **Hosts** section.

Adding a physical server

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

To add a physical server:

- 1 In the **My Inventory** view, click the **Physical** tab.
- 2 Click the **Add** icon, and then click **Physical Machine**.
The **Add Physical Machine Credentials** dialog box appears.
- 3 In the **DNS Name or IP** field, enter the FQDN or IP address of the server.
- 4 In the **User Name** field, enter an account for the server.
- 5 In the **User Password** field, enter the password for the preceding account.
- 6 If this server is new, or if you have removed the vRanger agent, select **Install agent on machine**.
In the **Agent Location** field, enter the preferred directory—on the physical machine—to which the physical client should be installed. The default is **C:\Program Files\Quest\vRangerPhysicalClient**.
- 7 In the **Port Number** field, 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 is 51000.
- 8 Click **Connect**.
The server appears in the My Inventory pane. You may also create a backup group to combine multiple physical servers into one backup job. For more information, see [Adding a custom backup group](#).

Installing the physical client manually

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

To install the physical client manually:

- 1 On the vRanger server, find the physical client installation file—**vRangerPhysicalClient.exe**.
The file is located in the **Service** subdirectory of the vRanger installation folder. The default is **C:\Program Files\Quest\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 following parameters:
 - **USERNAME:** Enter the name for an account with administrative **Log on as a service** rights on the source server.
 - **PASSWORD:** Enter the password for the preceding account.
 - **DOMAIN:** [Optional] If no value is entered, the preceding account is considered a local account. If you use a network account, enter the domain for the account here.

- **INSTALL_DIR:** [Optional] The default is **C:\Program Files\Quest\vRangerPhysicalClient**.
- **PORT:** [Optional] The default is 51000

Command examples

The following examples 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 In the My Repositories pane, right-click anywhere, and click **Add > Windows Share (CIFS)**.

i | **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](#).
 - 2 In the **Add Windows Network Share Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **User Name** and **Password:** Enter the credentials for accessing the CIFS share.
 - **Security Protocol:** Select a protocol, **NTLM** (default) or **NTLMv2**.

i | **IMPORTANT:** When adding a CIFS repository, the Security Protocol setting applies only to Virtual Appliance based operations.

For machine-based operations, the security protocol used is determined by the Windows configuration on the vRanger machine. By default, this is NTLM.

To use NTLMv2, you must edit the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\LmCompatibilityLevel** to set the **LmCompatibilityLevel** to 5.

 - **Server:** Enter the path to the preferred directory. Alternatively, you may enter the server IP or DNS name and click **Browse**, and browse to the correct path.
- !** | **CAUTION:** 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.
- 3 If you want these backups password-protected, select **Encrypt all backups to this repository**, enter a Password for the encrypted repository, and then confirm the password by re-entering it.

CAUTION: If you use the **Encrypt all backups to this repository** feature, make certain to retain the password you enter. There is no back-door or administrator-level password. If the password is unknown, the backups are not usable.

- 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 box.

Managing repository replication

If you are using Quest RDA or EMC® Data Domain® Boost (DD Boost™) repositories, you can also set up repository replication with another repository of the same type. For information on adding these types of repositories, see [Adding an EMC Data Domain Boost \(DD Boost\) repository](#) and [Adding a Quest Rapid Data Access \(RDA\) repository](#).

After the first repository is added to vRanger and another is available for replication, you can synchronize replication in one of three ways:

- Automatically, after a successful backup task to that repository.
- As a separate job on a scheduled basis.
- Manually, using 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 is also replicated to the replication repository.

To enable automatic savepoint replication:

- 1 On the Main toolbar, click **Tools > Options**.
- 2 Under the **Repositories** node, select **Replication**.

i **TIP:** You may also right-click the target repository in the **My Repositories** pane, and select **Repository Replication Options**.

- 3 Select **Enable savepoint replication for a successful backup job task**, and click **Ok**.

Scheduling repository synchronization

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

To schedule a repository synchronization:

- 1 On the Main toolbar, click **Tools > Options**.
- 2 Under the **Repositories** node, select **Replication**.
- 3 In the **Repository Replication Configuration** pane, find the **Repository Name** column, and select the applicable repository.
The repositories listed are the target repositories.
- 4 Select **Schedule repository synchronization**.
- 5 Configure the replication schedule as desired, using the following information as a guide.
 - a **Start:** In the drop-down list, select the time for the replication task to begin.

- b **Recurrence Pattern:** 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 value 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 value determines the month of the year in which the synchronization occurs.

x can be any value from 1 to 31. This value 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 value 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.

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

To synchronize a repository manually:

- 1 In the My Repositories pane, right-click the source or replication repository, and select **Synchronize**.
- 2 When the confirmation prompt appears, click **Yes**.

Disabling repository replication

To disable repository replication:

- 1 Do one of the following:
 - On the Main toolbar, click **Tools > Options**. Under the **Repositories** node, select **Replication**.
 - Right-click the replication repository, and select **Repository Replication Options**.

- 2 Clear the check box for **Enable savepoint replication for a successful backup job task** or **Schedule repository synchronization**, or both.
- 3 Click **OK**.


Replication is 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 In the **My Repositories** pane, right-click the replication repository, and click **Remove**.

 **IMPORTANT:** Ensure that **Delete all savepoints in this repository** is *not* selected.
- 2 Click **OK**.
- 3 In the **My Repositories** pane, click **Add**.
- 4 Select the appropriate repository type—either **EMC Data Domain Boost** or **Quest Rapid Data Access (RDA)**.
- 5 Complete the **Add Repository** dialog box as appropriate for your repository type, and click **OK**.
 Use the information in [Adding an EMC Data Domain Boost \(DD Boost\) repository](#) or [Adding a Quest Rapid Data Access \(RDA\) repository](#) as required.
- 6 When vRanger detects that the repository being added contains savepoint data and displays the **Warning: Existing Repository Found** dialog box, click **Import** to reconfigure the repository as a primary repository.

Configuring an SMTP server

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

To configure an SMTP server:

- 1 On the Main toolbar, click **Tools > Options**.
- 2 Under the **General** node on the **Configuration Options** page, click **Email Configuration**.
- 3 Populate the **SMTP Server** field.
- 4 Leave the default—25—in the **Port** field.
- 5 If your SMTP requires use of authentication or encryption, or both, click **Advanced**, and complete the following in the **SMTP Server Advanced** dialog box:
 - a To use authentication, select the **SMTP server requires authentication** check box, and enter the applicable user name and password. If applicable, include the domain followed by a backward slash before the name.
 - b To use SSL, select the **Connect using SSL** check box.
 - c Click **OK**.
- 6 In the **From Address** field, enter an email address from which notifications should be sent.
- 7 Click **Send Test Email**.
 The confirmation message confirms the server connection and successful transmission to the email address you entered.
- 8 Click **OK**.

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

Adding an email address

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

To add an email address:

- 1 On the Main toolbar, click **Tools > Options**.
- 2 Under the **General** node in 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 On the Main toolbar, click **Tools > Options**.
- 2 Under the **General** node, click **Email Configuration**.
- 3 In the **Email Addresses** section, select the email address that you want to update, and click **Edit**.
- 4 Update the information, and click **Update**.
The information is updated and moved to the **Email Addresses** section.
- 5 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 On the Main toolbar, click **Tools > 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 vRanger UI.

To configure display language:

- 1 On the Main toolbar, click **Tools > Options**.
- 2 Under the **General** node in the **Configuration Options** dialog box, click **Language**.

- 3 From the **Available Languages** drop-down list, 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 (VA)

vRanger uses a VA for Linux® FLR, standard backup and restore, and for replication to and from VMware® ESXi™ servers. The Virtual Appliance Configuration dialog box lets you add a configured VA to vRanger, and select the VA to use for Linux FLR. VAs deployed using the Virtual Appliance Deployment Wizard is added to vRanger automatically.

i | **NOTE:** You must have deployed and configured at least one VA before you can add it to vRanger. For instructions on VA deployment, see [Understanding the vRanger virtual appliance \(VA\)](#).

To configure a VA:

- 1 On the Main toolbar, click **Tools > Options**.
- 1 Under the **Virtual Appliance** node in the **Configuration Options** dialog box, click **Configuration**.
- 2 Select **Virtual Appliance Configuration**.
- 3 Click **Add**.

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

- 4 Browse the inventory tree, and select the VA.
- 5 Enter the root password for the VA, and click **OK**.
The VA 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 On the Main toolbar, click **Tools > Options**.
- 2 Under the **Licensing** node in the **Configuration Options** dialog box, click **Product**.
- 3 Click **Add New License**.
- 4 Browse to the location where you stored the license file, select it, and click **Open**.

The new license information is added to the dialog box.

- 5 Click **Apply**.
The dialog box 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 On the Main toolbar, click **Tools > Options**.
- 2 Under the **Licensing** node in the **Configuration Options** dialog box, click **Product**.
- 3 Click **Add New License**.
- 4 Browse to the location where you stored the license file, select it, and click **Open**.

The new license information is added to the dialog box.

- 5 Click **Apply**.

The dialog box refreshes with the new license information added. Under the **Currently Installed Licenses** heading, you see “License Type: Extended Trial,” and 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 them 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 **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 VM 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 in this dialog box.

i | **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 VM host licensing:

- 1 On the Main toolbar, click **Tools > 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.

The **Licensed Virtual Machine Hosts** section lists the managed hosts, the number of CPUs for each host, and the licensed features for that host. The **CPUs** column indicates the license “cost” of each host.

- 4 In the **Backup** or **Replication** column, or both, select or clear hosts until you have used all the available licenses, or until all 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 in this dialog box.

To manage physical machine licensing:

- 1 On the Main toolbar, click **Tools > 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.

The Licensed Physical Machine section lists the physical machines in the vRanger inventory, and indicates whether each machine is licensed for backup.
- 4 In the **Physical Backup** column, select or clear machines until you have used all the available licenses, or until all your machines are protected.
- 5 Click **OK**.

Managing job configuration settings

Through the Configuration Options dialog box, 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 running on the local vRanger machine.
per VA	The number of replication tasks running on a vRanger virtual appliance (VA).

To manage job configuration settings:

- 1 On the Main toolbar, click **Tools > 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 to 5): Default is 3.
 - On a host (1 to 4): Default is 1.
 - Per repository (≥ 1): Default is 3.
 - Locally (1 to 20): Default is 3.
 - Per VA (1 to 18): Default is 2.
- 4 In the **Timeout** section, populate the **Task Timeout** (1 - 999 hours) field—the default is 24.
- 5 In the **Minimum space needed on host (GB)** field, enter a number—the default is 6.
- 6 Click **OK**.

Enabling cataloging options

The vRanger cataloging function must be enabled before data collection can begin. The Catalog dialog box 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 On the Main toolbar, click **Tools > 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 FLR 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 following directories. Path filtering is determined by entries in the **PathFilterTokens.txt** file, at **C:\Program Files\Quest\VRanger\Configuration**.

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

- **File:** By default, vRanger does not catalog files of the following type. File filtering is determined by entries in the **FilesFilterTokens.txt** file, at **C:\Program Files\Quest\VRanger\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 occur 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. For VMware® ESXi™ hosts, configure the transport order to **VA-based HotAdd > VA-based Network**.

To manage transport failover settings:

- 1 On the Main toolbar, click **Tools > 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 select it, and use the up or down arrows to move the item higher or lower in the order.

Configuring the number of times a backup process tries to run

By default, a backup job that fails tries to run again. The default for the maximum number of times that a job is automatically run again is three. If vRanger appears to run a failing job repeatedly, check the configuration file to ensure that the default is three. Configuration information is maintained in the configuration file for the Quest vRanger Service, **Vizioncore.vRanger.Service.exe.config**. This file is located in the vRanger installation directory. By default this location is: **C:\Program Files\Quest\vRanger**.

CAUTION: Open and update this file only when necessary. Change only the values listed in the following procedure unless directed otherwise by Quest Support.

To open the file and view the applicable information:

- 1 Locate the **Vizioncore.vRanger.Service.exe.config** file, and use a text editor to open it.
There are two keys that limit the number of retries:

```
<add key="BackupRestoreMaxTries" value="3"/>
```

```
<add key="BackupMaxTriesOnVMotion" value="3"/>
```
- 2 Verify that these keys are set to **3**.
 - If so, close the file without changes.
 - If the value is different or blank, for example, **value=""**, update the entry, save the file, and restart the Quest vRanger Service. For more information on restarting the service, see [Restarting the Quest vRanger Service](#).

Supplemental instructions: additional repository types

The main procedures describe the process for adding a CIFS repository. To add one of the other types, see the applicable topic:

- [Adding an NFS repository](#)
- [Adding a NetVault SmartDisk repository](#)
- [Adding an EMC Data Domain Boost \(DD Boost\) repository](#)
- [Adding a Quest Rapid Data Access \(RDA\) repository](#)

i | **NOTE:** Quest RDA and EMC DDB repositories support repository replication. For more information on configuring repository replication, see [Managing repository replication](#).

Adding an NFS repository

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—the vRanger machine and any VMware® ESXi™ 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 entry would look like:

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

To add an NFS repository:

- 1 In the My Repositories pane, right-click anywhere, and click **Add > NFS**.
- 2 In the **Add Network File Share Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter 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 subdirectory in the export directory.
 - **Target Directory:** Enter a subdirectory of the NFS Export directory. This subdirectory is the location to which savepoints are written.

i | **CAUTION:** If you use the **Encrypt all backups to this repository** feature, make certain to retain the password you enter. There is no back-door or administrator-level password. If the password is unknown, the backups are not usable.

- 3 If you want these backups password-protected, select **Encrypt all backups to this repository**, enter a Password for the encrypted repository, and then confirm the password by re-entering it.
- 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 box.

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 5 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
 - **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.

Adding a NetVault SmartDisk repository

The following process describes how to add a Quest NetVault SmartDisk repository.

To add a NetVault SmartDisk repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > NetVault SmartDisk (NVSD)**.

- 2 In the **Add NVSD Repository** dialog box, complete the following fields:

- **Repository Name:** Enter a descriptive name for the repository.
- **Description:** [Optional] Enter a long-form description for the repository.
- **DNS Name or IP:** Enter the IP or FQDN for the NetVault SmartDisk server.
- **User name and Password:** Enter the credentials for the account that vRanger should use to connect to the repository.

i | NOTE: If you have enabled WebDav authentication on your NetVault SmartDisk server, use those credentials here. If you have not enabled WebDav authentication, no credentials are required. For more information, see the *Quest NetVault SmartDisk Administrator's Guide*.

- **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 are available in a drop-down menu.

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

- **Port:** Enter the port that vRanger should use to connect to the NetVault SmartDisk server—the default is 37453.
- **Timeout:** Enter a value that determines how long to wait before terminating an idle connection.

i | NOTE: Encryption is not supported for NetVault SmartDisk repositories.

- 3 Click **OK**.

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

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 4 If vRanger finds existing savepoints, click the applicable button:

- **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
- **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
- **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats 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 In the **My Repositories** pane, right-click the NetVault SmartDisk repository, and select **Remove**.

The **Remove Repository** dialog box 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.

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

CAUTION: This step deletes the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.

- 3 If you chose to delete the savepoints, click **OK** when the **Delete Savepoints** dialog box appears.
The savepoints are deleted, along with the repository.

Adding an EMC Data Domain Boost (DD Boost) repository

The following process describes how to add an EMC® Data Domain® Boost (DD Boost™) repository. The following procedure assumes that:

- You have a properly configured Data Domain appliance that is accessible to the vRanger machine.
- DD Boost is licensed and enabled on the Data Domain device.
- You have designated a DD Boost User account.

NOTE: For information on setting up the preceding configurations, see your EMC Data Domain documentation.

To add an EMC Data Domain Boost Repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > EMC Data Domain Boost (DDB)**.
- 2 In the **Add EMC Data Domain Boost Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter the IP or FQDN for the Data Domain device.
 - **User name and Password:** Enter the credentials for the account selected as the **DD Boost User** on the Data Domain device.

NOTE: For information on creating a DD Boost user name, see your Data Domain documentation.

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

- 3 Click **OK**.

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

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 4 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.

Configuring repository replication

Replication is configured through the My Repositories pane of the vRanger interface. The following procedure assumes that:

- 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. You do not have to add this device to vRanger.
- All devices used in replication have a Replicator license installed.

To configure replication:

- 1 In the **My Repositories** pane, right-click the Data Domain Boost repository, and select **Configure Replication**.
- 2 In the **Configure Replication** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter the IP or FQDN for the Data Domain device.
 - **DD Boost User name and Password:** Enter the credentials for the account selected as the **DD Boost User** on the Data Domain device.
 - **Storage Unit:** This field specifies 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 creates 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 box.

After replication is configured for a repository, the **Configure Replication** option is disabled for that repository.

i | IMPORTANT: After a repository is configured for replication, you must select a synchronization method before replication occurs. For information on manual and automatic synchronization options, see [Managing repository replication](#).

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 box to view free space for the repository.

- 1 In the **My Repositories** pane, right-click the **EMC Data Domain Boost** repository, and select **Edit Replication Configuration**.
- 2 In the **Edit EMC Data Domain Boost Repository** dialog box, edit any of the following fields:
 - **Repository Name**
 - **User Name**
 - **Password**
- 3 Alternatively, view the **Free Space** field for up-to-date information about this repository.
- 4 If you edited any of the fields, click **Update**.

i | NOTE: For information on disabling repository replication, see [Disabling repository replication](#).

Removing an EMC DD Boost repository

The process for removing an EMC® Data Domain® Boost (DD 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 In the **My Repositories** pane, right-click the applicable EMC Data Domain Boost repository, and select **Remove**.

The **Remove Repository** dialog box 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. To remove the storage unit associated with the repository, you need to remove the savepoints. If replication is configured for this repository, you are also given the option to delete the savepoints in the replicated repository.

- 2 Select **Delete all savepoints in this repository**.

CAUTION: This step deletes the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.

- 3 To delete the savepoints in a replicated repository, select **Delete all savepoints in replication repository**.
- 4 Click **OK**.
- 5 When the **Delete Savepoints** dialog box appears, select **Delete DD Boost storage unit**, and click **OK**.

The savepoints are deleted, along with the storage unit associated with the repository.

Adding a Quest Rapid Data Access (RDA) repository

The Quest 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 Quest DR Series include the following appliance types:

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

The Quest DR Series supports Common Internet File System (CIFS), Network File System (NFS), and RDA protocols. The RDA protocol provides a logical disk interface for the Quest DR Series system, and enables better coordination and integration between vRanger and the Quest DR Series system and provides for client-side deduplication of vRanger backups.

NOTE: Quest recommends that you use the RDA protocol when using a Quest DR Series system as a repository.

For more information about the Quest DR Series systems, see the *Quest DR Series System Administrator Guide*.

Adding a Quest DR Series system as an RDA repository

The following process describes how to add a Quest Rapid Data Access (RDA) repository. The following procedure assumes that:

- You have properly configured the Quest DR Series appliance that is accessible to the vRanger machine.
- You 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
- You have designated an RDA User account.
- You configured your firewall to enable the following TCP ports:
 - 9904
 - 9911
 - 9915
 - 9916
 - 9920

i | NOTE: For information on setting up the preceding configurations, see your Quest DR Series documentation.

To add a Quest DR Series system as an RDA repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > Quest Rapid Data Access (RDA)**.
- 2 In the **Add Quest Rapid Access Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter the DNS name or IP address of the Quest DR Series system.
 - **RDA Username:** Enter a user account that can be used to log in to the device. On the Quest 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 an account or delete the existing account.
 - **RDA Password:** Enter the password for the user account. The default is: **St0r@ge!**
 - **Logical Storage Unit:** Enter the name of the storage container. Ensure that the container is created before you add the device. You cannot add the device 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.
- 3 Click **OK**.

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

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.
- 4 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.

- **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
- **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.

Configuring repository replication

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

i | TIP: For more information on repository replication, see the *Quest DR Series System Administrator Guide*.

The following procedure assumes that:

- You have configured at least one Quest RDA repository in vRanger.
- At least one other Quest DR Series appliance with RDA is available in your environment to serve as the replication target. You do not have to add this device to vRanger.

To configure Replication:

- 1 In the **My Repositories** pane, right-click the Quest Rapid Data Access (RDA) repository, and select **Configure Replication**.
- 2 In the **Configure Replication** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** The DNS name or IP address of the Quest DR Series system.
 - **RDA Username:** Enter a user account that can be used to log in to the device. On the Quest 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 an account or delete the existing account.
 - **RDA Password:** Enter 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. You cannot add the device 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.
- 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 box.

After replication is configured for a repository, the **Configure Replication** option is disabled for that repository.

i | IMPORTANT: After a repository is configured for replication, you must select a synchronization method before replication occurs. For information on manual and automatic synchronization options, see [Managing repository replication](#).

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 box to view free space for the repository.

- 1 In the **My Repositories** pane, right-click the Quest Rapid Data Access (RDA) repository, and select **Edit Replication Configuration**.

- 2 In the **Edit Quest Rapid Data Access (RDA)** dialog box, edit any of the following fields:
 - **Repository Name**
 - **User Name**
 - **Password**
- 3 Alternatively, view the **Free Space** field for up-to-date information about this repository.
- 4 If you edited any of the fields, click **Update**.

i | **NOTE:** For information on disabling repository replication, see [Disabling repository replication](#).

Removing a Quest Rapid Data Access (RDA) repository

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

To remove a Quest Rapid Data Access (RDA) repository:

- 1 In the **My Repositories** pane, right-click the Quest Rapid Data Access (RDA) repository, and select **Remove**.

The **Remove Repository** dialog box 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. To remove the storage unit associated with the repository, you need to remove the savepoints. If replication is configured for this repository, you are also given the option to delete the savepoints in the replicated repository.

- 2 Select **Delete all savepoints in this repository**.

i | **CAUTION:** This step deletes the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.

- 3 To delete the savepoints in a replicated repository, select **Delete all savepoints in replication repository**.
- 4 Click **OK**.
- 5 When the **Delete Savepoints** dialog box appears, click **OK** to delete the savepoints.

Adding a Quest DR Series system as a CIFS repository

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

To add a Quest DR Series system as a CIFS repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > Windows Share (CIFS)**.
- 2 In the **Add Windows Network Share Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **User Name** and **Password:** Enter the credentials for accessing the CIFS share.
 - **Security Protocol:** Select a protocol, **NTLM** (default) or **NTLMv2**.
 - **Server:** Enter the UNC path to the desired repository directory. Alternatively, you may enter a partial path and click **Browse** to find the target directory.

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

i | **IMPORTANT:** Do *not* select **Encrypt all backups to this repository**. Using encryption or compression with deduplicated repositories limits or disables deduplication. Encryption and compression should not be used with any repository type that provides deduplication.

- 3 Click **OK**.

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

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 4 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
 - **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.
- 5 Click **Next**.

Adding a Quest DR Series system as an NFS repository

i | **NOTE:** Adding a Quest DR Series system as a Network File System (NFS) repository does not take advantage of the full function of the Quest DR Series system. Quest recommends adding the Quest DR Series as an RDA repository.

To add a Quest DR Series system as an NFS repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > NFS**.
- 2 In the **Add Network File Share Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter the IP or FQDN for the repository.
 - **Export Directory:** Specify the export directory, which is similar in concept to a network share. You must create a target subdirectory in the export directory.
 - **Target Directory:** Specify a subdirectory of the NFS export directory. This directory is the location to which savepoints are written.

i | **IMPORTANT:** Do *not* select **Encrypt all backups to this repository**. Using encryption or compression with deduplicated repositories limits or disables deduplication. Encryption and compression should not be used with any repository type that provides deduplication.

- 3 Click **OK**.

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

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 4 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.

- **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.

Understanding the vRanger virtual appliance (VA)

The vRanger VA is a small, pre-packaged Linux[®] distribution that serves as a platform for vRanger operations away from the vRanger server. With VAs, the workload can be spread across the other CPUs available to a host. This feature provides increased reliability and scalability over operations.

vRanger uses the VA for the following VMware[®] functions:

- Replication to and from VMware[®] ESXi[™] hosts.
- File-level restore (FLR) from Linux machines.
- Optionally, for backups and restores.

i | NOTE: The VA is not designed to be used with Hyper-V[®] machines.

The VA 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 choose a host in the cluster to which the VA should be associated.

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

When configuring the VA, consider 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 VA, carefully consider an appropriate size for the VA scratch disk. For more information, see [The VA scratch disk](#).

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

- [The VA scratch disk](#)
- [Deploying the VA by using the Virtual Appliance Deployment Wizard](#)
- [About deploying multiple VAs](#)
- [Changing the VA configuration](#)
- [Configuring VA networking](#)
- [Disabling Distributed Resource Scheduler \(DRS\) for the vRanger VA](#)

The VA scratch disk

If you want to perform replication tasks with your VA, 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 file contains 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 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, by using 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, the original destination disk blocks are read from the undo file and written to the destination disk to roll back the failed replication. This process 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. 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 MB, 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 following topics to guide you in determining the proper size for the scratch disk.

- [Strategies for sizing the scratch disk](#)
- [Options for a smaller scratch disk](#)
- [The scratch disk on the source host](#)
- [Scratch disk location](#)

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 are replicating, their total disk size, replication frequency, and the data change rate per VM. It is important to understand all 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 size the scratch disk properly, without wasting storage space, is to use historical replication data. This data is available in the Replicate Task Reports, in the vRanger My Reports view, for the applicable VMs. 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, Quest 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 following numbers are examples given to illustrate the concept:

For example, if you have four VMs 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
3	100 GB	20%	20 GB
4	60 GB	5%	3 GB

For the preceding VMs, you would need approximately 48 GB of disk space for the undo files, plus a buffer of approximately 10%, for safety's sake. In the example, an appropriate estimate for the scratch disk size for the preceding VMs 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.

CAUTION: If the scratch disk runs out of space, replication tasks fail with a “FATAL cant_write can't write vix” error.

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 case you need to fail over to the disaster recovery (DR) site, the replication job reverses direction and starts replicating changes back to the product site—the original source host or cluster. For this process 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 that you place the disk on a datastore with block sizes large enough to support the expected VMDK. The following list 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 VA by using the Virtual Appliance Deployment Wizard

vRanger uses a virtual appliance (VA) for replication to and from VMware® ESXi™ hosts, for FLR from Linux® machines, and optionally for backups and restores. The VA Deployment Wizard offers a simple method to deploy VAs one at a time. You can run the wizard for each VA that you need to deploy, or alternatively deploy multiple VAs at once as described in [About deploying multiple VAs](#).

NOTE: The automated deployment process requires DHCP. To deploy a VA through the Virtual Appliance Deployment Wizard, you must have a DHCP server accessible to the VA during deployment.

To deploy the VA using the wizard, complete the following tasks:

- [Starting the Virtual Appliance Deployment Wizard](#)
- [Selecting a VA host](#)
- [Selecting VA deployment options](#)
- [Configuring the VA IP address](#)
- [Confirming VA deployment selections](#)

Starting the Virtual Appliance Deployment Wizard

Complete the steps in the following procedure to start the wizard.

To start the wizard:

- 1 Do one of the following:
 - In the **Tools** menu, click **Virtual Appliance Deployment Wizard**.
 - In the **My Inventory** view, right-click 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 box, click **Deploy Virtual Appliance**.
- 2 When the **Virtual Appliance Deployment Wizard** appears, 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 VA 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 In the vCenter view in the center of the **Host Selection** dialog, select the host to which the VA should be deployed.
- 2 If you select a single host (a host that is not part of a cluster), select **Associate VA with specific host**.
OR
If you select a host on a cluster, select either:
 - **Associate VA with specific host** - this option results in the VA being used only for VMs on the specified host.
 - **Associate VA with cluster** - this option designates the VA as being shared among all of the hosts on the cluster. If a host in the cluster does not have a VA associated locally, then a shared VA will be used.

- 3 Click **Next**.

Selecting a VM Folder

Select a datacenter or VM folder for the virtual appliance.

- 1 Click on the desired datacenter or VM folder.
- 2 Click **Next**.

Select a Computing Resource

In this dialog, you will select a computing resource for the VA. A computing resource provides the CPU and memory used by the virtual appliance.

- 1 From the vCenter view in the center of the dialog, select a computing resource. Select either a host, vApp, or resource pool.

Selecting VA deployment options

In the VA Deployment Options dialog box, you can configure the VA's name and allocated resources. Also, you can configure the size of the VA's scratch disk. For more information, see [The VA scratch disk](#).

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 VA.

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 entry is sufficient for two concurrent tasks per VA.
- **Quest Recommended:** This setting allocates two CPUs and 1 GB of RAM. This entry is sufficient for four concurrent tasks per VA.
- **Quest Recommended (with RDA repository):** This setting allocates four CPUs and 2 GB of RAM. This entry 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** field, select the datastore for the VA's primary disk.
- 4 In the **VM Storage Policy** field, select the desired storage policy for the VA. Compatible storage policies for the selected Datastore are presented. The default selection is **Datastore Default**.

i **NOTE:** Once applied to a vRanger VA, a VM Storage Policy cannot be changed through vRanger. Editing the Storage Policy must be done through the VMware vSphere client.

- 5 In the **Network Assignment** field, select the network for the VA's primary NIC.
- 6 If you want to be able to perform replication tasks on this VA, configure a scratch disk on the VA—for more information, see [The VA scratch disk](#).

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.

i **NOTE:** If you are upgrading an existing VA and want to migrate the scratch disk, do not select this option. For more information, see the topic on upgrading the vRanger VA in the *Quest vRanger Installation/Upgrade Guide*.

- 7 Do one of the following:

- In the **VA Password** and **Confirm password** field, enter a new password for the VA. If you change the password, this password becomes the default for subsequent VA deployments performed during this session.
- Leave the default password of **vzroot1**.

i | **NOTE:** You may change the password in the future as described in [Changing the VA configuration](#).

- 8 If you want to perform FLR from backups of Linux® VMs, configure a VA to use for Linux FLR.

Select **Use virtual appliance for Linux File Level Restore**.

If you select this option for a different VA in the future, that VA is used for Linux FLR.

Configuring the VA IP address

On the VA IP Address Configuration page, you must configure the network configuration for the VA's primary network interface card (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.

i | **NOTE:** If you need to change your network configuration after the VA is deployed, you need to do so from within the VA. For instructions, see [Configuring VA networking](#).
- 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 change to the information.
- 2 [Optional] 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 VA to multiple hosts, it may be helpful to create a VM template from the configured VA.

To create a template:

- 1 From the VI Client, right-click the configured VA, select **Template**, and then click **Clone to Template**.
- 2 Enter the template name and select a location, and click **Next**.
- 3 Select the host or cluster on which the template is to be stored, and click **Next**.

- 4 Select the datastore in which the template is to be stored, and click **Next**.
- 5 In the **Disk Format** dialog box, select **Same format as source**, and click **Next**.
- 6 Review the configurations, and click **Finish**.

About deploying multiple VAs

The VA must be deployed to any VMware® 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 VA requires that if a VA is used on one host or cluster in a replication job, a VA must be used on both the source and destination host or cluster. In other words, VAs, when used for replication, must be used in pairs.

For a few VAs, using the vRanger UI to deploy the VAs is sufficient. To streamline the process of deploying a high number of VAs, Quest recommends that you deploy and configure the VA once, and save it as a template to be used for additional deployments using the VMware vSphere® PowerCLI™. The [Creating a template](#) procedures include an optional step for creating a template from the configured VA.

Changing the VA configuration

After a VA is deployed, you may change its configuration from the vRanger UI.

IMPORTANT: Reconfiguring the VA might cause it to be power cycled, possibly affecting configured jobs. You are prompted with the jobs configured for the VA; review these entries carefully to avoid service interruption.

To change the VA configuration:

- 1 Do one of the following:
 - In the **My Inventory** view, right-click the VA you want to change, and then click **Virtual Appliance Configuration**.
 - On the Main toolbar, do the following:
 - a Click **Tools**, and then click **Options**.
 - b In the **Configuration Options** dialog box, select **Virtual Appliance**.
 - c Under **Configure Existing Virtual Appliances**, select the VA to change, and then click **Edit**.

The **Modify Virtual Appliance Configuration** dialog box appears.

- 2 In the **Virtual Appliance Properties** section, do the following:
 - a Next to **VA Name**, confirm the name of the VA.
 - b Next to **IP Address**, confirm the IP vRanger uses to connect to the VA.

By default, vRanger connects to the vRanger VA using the first IP address reported by the VMware vSphere® API. This IP is displayed under Virtual Appliance Properties in the IP Address field. If you have only one NIC configured on the VA, continue to the next step.

If you have more than one NIC configured on the VA, 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 VA, and want vRanger to connect to the second NIC, select **Override IP Address** to enable the **IP Address** field, and then enter the IP for the second NIC in the IP Address field.

i | **NOTE:** For information about manual network configuration procedures, see [Configuring VA networking](#).

- 3 In the **Virtual Appliance Options** section, use the **Select an option** drop-down list to change the resources allocated to the VA as needed.

You may select from one of the following options:

- **Minimum Requirement:** This setting allocates one CPU and 512 MB of RAM. This entry is sufficient for two concurrent tasks per VA.
- **Quest Recommended:** This setting allocates two CPUs and 1 GB of RAM. This entry 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 list.

i | **NOTE:** If you are using a VMware® vCenter™ version lower than 6.0, the ability to change the datastore of the VA scratch disk used for replication is not available.

- 5 If you want to perform replication tasks on this VA, configure a scratch disk on the VA—for more information, see [The VA scratch disk](#).

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](#).

- 6 [Optional] Enter a new password for the VA in the **VA Password** field, and then re-enter it in the **Confirm password** field.

- 7 If you want to perform FLR from backups of Linux® VMs, select **Use virtual appliance for Linux File Level Restore** to configure a VA to use for Linux FLR.

If you select this option for a different VA in the future, that VA is used for Linux FLR.

Configuring VA networking

If you have deployed one or more VAs using a static IP configuration, and you need to change that configuration after deployment, you must do so from within the VA.

To configure VA networking, complete the following tasks:

- [Logging in to the VA](#)
- [Changing the hostname](#)
- [Changing the network configuration](#)
- [Configuring DNS](#)

Logging in to the VA

Complete the steps in the following procedure to log in to the VA.

To log in to the VA:

- 1 From your VI Client, find the VA, and power it on.
- 2 Using the console, log in as **root** to the VA, using the following default credentials:
 - 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 Log in to the VA.
- 2 Using the VI text editor, edit the network configuration file:
`vi /etc/sysconfig/network`
- 3 Change the hostname to the new value, and 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 Log in to the VA.
- 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 following text is shown as an example:

```
ONBOOT=yes
SERVICE=ipv4-static
IFACE=eth0
IP=<static IP of VA>
GATEWAY=<IP of gateway>
PREFIX=<this entry 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, you might have to configure the VA with DNS nameserver information. The steps to change an existing configuration are the same.

To configure DNS:

- 1 Log in to the VA.
- 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 Distributed Resource Scheduler (DRS) for the vRanger VA

If the vRanger appliance is configured on hosts that are part of a DRS-enabled cluster, the VAs might be migrated to a different host due to DRS and high-availability (HA) activity. This change causes replication failures as the VA is no longer in the expected location—as configured in vRanger. To prevent this disruption, it is important to disable DRS for the vRanger VAs.

To disable DRS for the vRanger VA:

- 1 In your VI client, find the cluster or clusters that contain the vRanger VAs.
- 2 Right-click the cluster, and select **Edit Settings**.
- 3 In the **Cluster Settings** dialog box, select **Virtual Machine Options**.
- 4 Search for or find the vRanger VAs. Right-click each VA, and select **manual**.

Using vRanger

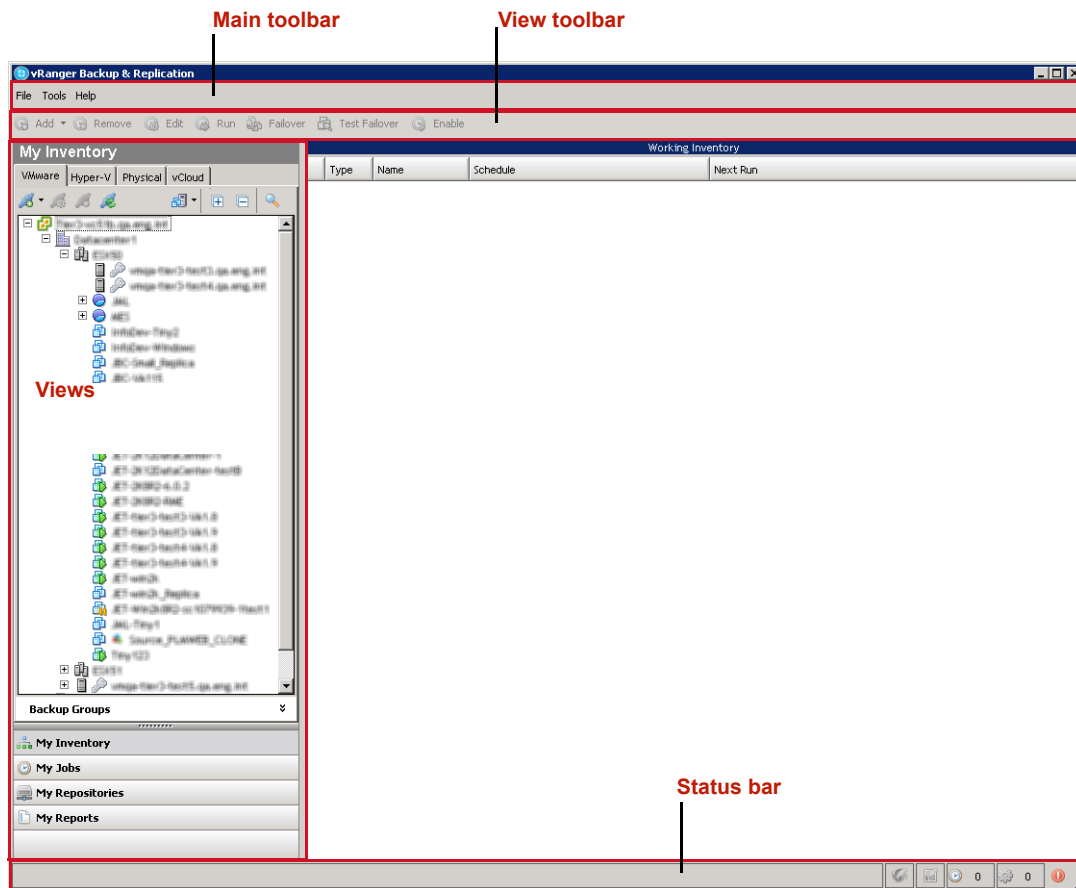
- [vRanger user interface \(UI\) overview](#)
- [My Inventory view](#)
- [My Jobs view](#)
- [My Repositories view](#)
- [My Reports view](#)

vRanger user interface (UI) overview

The vRanger UI consists of four views: My Inventory, My Repositories, My Jobs, and My Reports. Each view features the Main toolbar and each pane's toolbar.

In the left pane within these views, data appears nested in a tree structure. In general, the information that appears 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 UI 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 are looking for. Restore operations can be started from the items in the results pane.
- **Startup Wizard:** Select this command to start the Startup Wizard, through which you can add a VC, host, repository; license a host; set up a mail server; and add email addresses.
- **Virtual Appliance Deployment Wizard:** Select this command to start the Virtual Appliance Deployment Wizard, through which you can deploy and configure vRanger virtual appliances for backup and replication.
- **Options:** When you select this command, the Configuration Options dialog box appears. 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 box 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 box, you can import a license file.
- **What's New:** Select this command to view a list of new features and enhancements in your version of vRanger.
- **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 box, which includes hardware resources, components, software environment, and internet settings. The dialog box 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 Quest 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 following table, 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 this icon to enable log files to help in troubleshooting Windows® vRanger service.



Client Support Logging: Right-click this icon to enable log files to help in troubleshooting the vRanger UI.



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 appears 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 My Inventory pane on the left and the Working Inventory pane on the right. This view is the one from which most backup and replication jobs are added.



My Inventory pane

The My Inventory pane offers four tabs—one to show the infrastructure of VMware® VMs, one to show the inventory of Hyper-V® VMs, one for physical servers, and one to show the inventory of VMware vCloud Director® servers.

From this pane, you can navigate and refresh the tree; add or remove connections to VCs or VMware® ESXi™ servers; and add credentials to virtual servers. You can also add a backup or replication job.

My Inventory pane toolbar and icons

The toolbar is located under the Main 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 appears next 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 appears next to the listing.




Remove: When you click to select 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 option reloads the window and displays the job in **My Jobs: Current Jobs**.

The following icons 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 on the node that you have selected in the **My Inventory** pane.

- **VMware tab:** Click to add a VMware® vCenter™, VMware® 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 VMware® 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 ESXi 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 appears (VC). A Loading Inventory message appears while the pane refreshes.



Inventory: Click one of these icons to display the **Hosts and Clusters** or **Virtual Machines and Templates** view.

NOTE: This option applies only to the VMware tab.



Expand All: This option 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 appears in the pane.




Search: Click this icon to display the Search Inventory field and find a specific item in the selected inventory pane.

NOTE: This option applies only to the VMware and vCloud tabs.

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 following icons 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 a VMware® ESXi™ host.



ESX(i) Host Disconnected Right click this icon to connect to the host.



ESX(i) Host Error: Right-click this icon to determine the error.



ESX(i) Host Maintenance Mode: Indicates that the host is in maintenance mode, and might not be available.



Hosts and Clusters



vApp: VMware vSphere® vApp(s)™ can include any applications running on any operating system, 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 Changed Block Tracking (CBT) enabled. Click the icon to disable.



Virtual Machine Encrypted: This icon indicates that the VM is encrypted.



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 Inventory view

When the Virtual Machines and Templates Inventory option is selected, the following icons 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 vCenter.



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 this icon to display the VMs in this folder.



Virtual Machine Folder Open: Clicking this icon collapses the view and closes the folder.



Virtual Machine Encrypted: This icon indicates that the VM is encrypted.



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 Service: Click this icon to expand the list of services.



Running Virtual Machine: 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 the preferred VM, and click this icon to start a backup job for this VM.



Replicate Virtual Machine: Right-click the preferred VM, and click this icon to start a replication job for this VM.



Change Block Tracking: This icon indicates that a job has CBT 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.



Data Center: Selecting this icon expands the data center and displays all objects in it.



Disconnected: This icon indicates that 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: VMware vSphere® vApp(s)™ can include any applications running on any operating system, 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 that the VC is licensed for both backup *and* replication.



Backup or Replication License: Indicates that 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 the lack of communication include 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

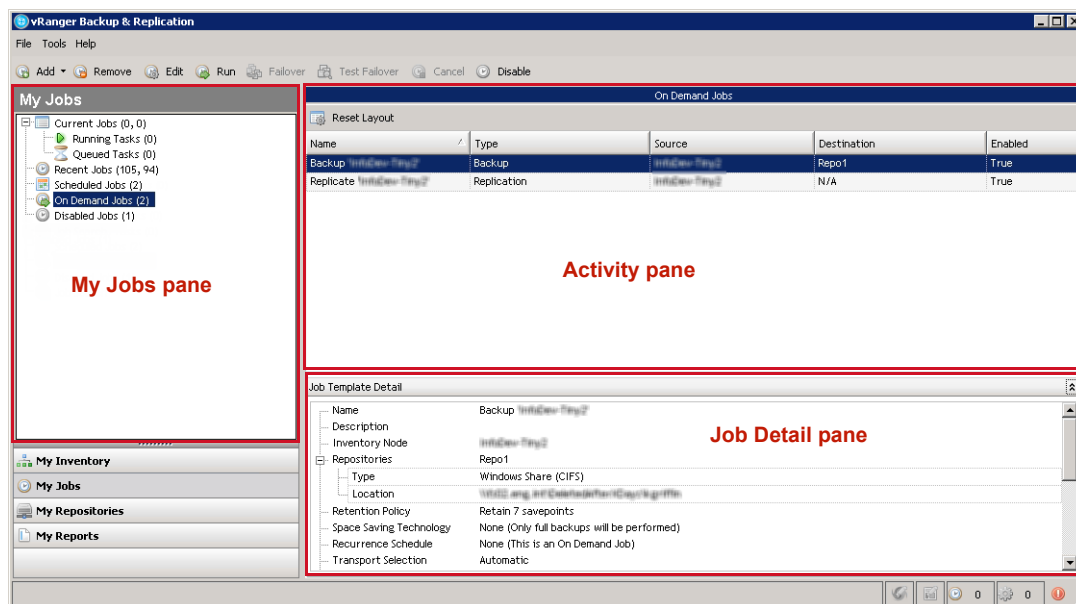
The Working Inventory pane lists any configured jobs for the node selected in the Inventory pane.

The data presented here—including Type, Name, and Schedule—can be sorted by dragging a column heading. When you click, an arrow is displayed 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 the headings. Taking this step groups current jobs by that data type.

If you click to select 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—selected 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 selected in the **Activity** pane.



Edit: With a job selected, click 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 appears 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 appears next to the listing.



Run: Click this command to run the job that is selected. This option 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: Indicates 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. Selecting this node also displays the **Recent Job Filters** section, where you can set up specific search criteria. For more information, see [Filtering job searches](#).



Scheduled Jobs: When this node is selected, the **Activity** pane is divided into two tabs: **List View** and **Calendar View**. The items listed on both views are those jobs and tasks 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 might have been run, but are currently disabled. From this view, any disabled job can be enabled.

Filtering job searches

The **Recent Job Filters** section contains several subsections that you can use to filter the jobs that are displayed in the **Recent Jobs** pane.

To filter a job search:

- 1 In the **My Jobs** pane, click **Recent Jobs**.
- 2 In the **Search Options** section, select the applicable options:
 - **Rows/Page:** Enter the number of rows to display on each page in the Recent Jobs pane.
 - **Order By:** Indicate whether the list should sort by **Start Date**, **End Date**, **Job Name**, or **Job Status**. Also indicate whether to display the list in an **Ascending** or **Descending** order.
 - **Save as default search:** If you want to save the preceding search criteria as your default, select the check box.
- 3 In the **Job Filter** section, select the applicable options:
 - **Job Type:** Select **All Jobs**, **Backup Jobs**, **Replication Jobs**, **Restore Jobs**, or **Repository Replication**.
 - **Job Status:** Select **All**, **Succeeded**, **Failed**, **Aborted**, or **Canceled**.
 - **Date Range** or **Days in the past:** To specify a range of dates, set a date in **Start Time** and **End Time**. To specify a specific number of days instead, enter a number in the box.
- 4 In the **Job Task Filter** section, select the applicable options:
 - **Task Status:** Select **All**, **Succeeded**, **Failed**, **Aborted**, or **Canceled**.
 - **Repository:** Select **All** or select a specific repository.
 - **Host:** Select **All** or select a specific host.
 - **Machine name:** To view the tasks for a specific system, enter the name of the machine.
- 5 Click **Search**.

The information displayed in the Recent Jobs pane updates to reflect your selections. You can select any item in the list and view additional information about the job or task on the **Detail** and **Log** tabs in the **Job Detail** pane.

Activity pane

The data that appears in the Activity pane depends on 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:** Displays the name of the job.
- **Type:** Refers to job type.
- **Source:** Indicates the host or cluster in which the source server resides.
- **Destination:** For backup jobs, this category indicates the repository selected in the backup job. For restore and replication jobs, this category 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:** Indicates length of time that the job or task took to complete.
- **Message:** Displays more information about a task or job. For example, "Host is not licensed" or "Please check host credentials."

If you select the Current Jobs node, all the preceding categories display in the Activity pane. In addition, data for the following categories is included: Throughput and Progress.

In the vRanger UI, you can clearly distinguish the running job (parent) and the running task (child). To view details for either job or task, click to select the object. Either the Job Detail or Task Detail area appears 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 a job or task in the Activity pane, a Go To menu appears. The available commands list depends on 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 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 area, 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 select any entry and additional information about the job appears in the Job Template Detail area below the results, which you can save.

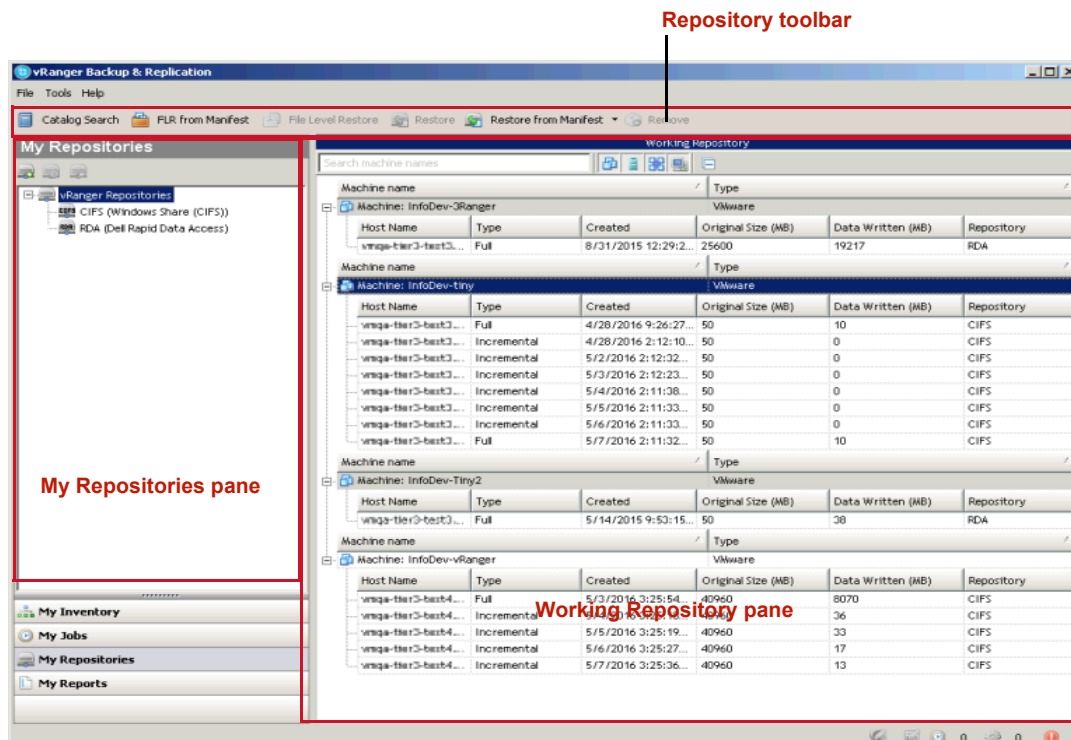
Type	<ul style="list-style-type: none"> • Historic • On Demand • Scheduled • Disabled
Task Status	<ul style="list-style-type: none"> • All • Success • Failed • Aborted • Canceled
Repository	The values included here depend on the repositories that were added previously.
Job Type	<ul style="list-style-type: none"> • All • Backup • Restore • Replication
VM Name	To search by VM name, populate this field.
Job Status	<ul style="list-style-type: none"> • All • Success • Failed • Aborted • Canceled
To Date	The default is (none). When you click the drop-down list, a calendar view appears.
Host or VC	The default is (All). This list is populated with the hosts and VCs that were added previously.

Job Detail area

The **Job Detail** area contains job configuration data and 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 in 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 view is the one 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 Restore** dialog box.



FLR From Manifest: When you want to restore an individual file from the Manifest, click this icon.



Restore: When you click to select 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 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 box. 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 filtered by inventory type and/or search criteria. Selecting the top-level **vRanger Repositories** node will display savepoint contents for all repositories. Selecting a specific repository displays only the contents of that repository.

- The data in this pane is organized under these columns: Machine Name
- Host Name
- Type—Full, Differential, Incremental
- Start Time
- Original Size (MB)
- Data Written (MB)
- Repository

Grouping

The Working Repository pane groups savepoints by machine name and can be filtered by machine type or to match search results.



Search: Free-form search that filters displayed machine names based on characters entered into the search form.



Machine type: Toggles the display of savepoints based on machine type. Machine types can be selected in any combination. Options are:

- VMware
- Hyper-V
- vCloud
- Physical



Expand All: Expand all machine entries to view savepoints.



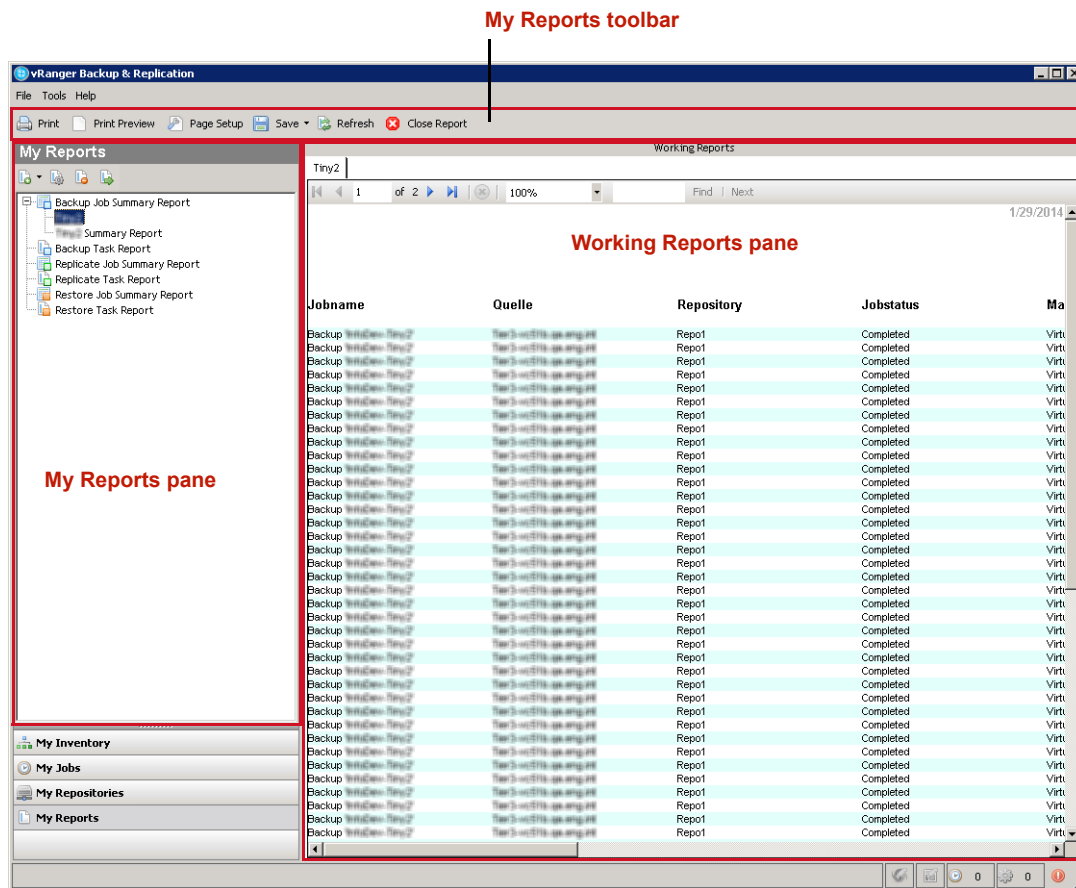
Collapse All: Collapse all machine entries to hide savepoints.

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 is displayed 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 and email them. The My Reports view is made up of three areas: My Reports on the left, which lists all reports for backups and restores; Working Reports on the right, which displays the report itself; and the Report Template Detail area on the lower right.



My Reports toolbar

The toolbar is located under the Main 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.

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.



Print: Use this icon to print the report displayed in the **Working Reports** pane.



Print Preview: Use this icon to review the report before printing it.



Page Setup: Use this icon to adjust the report before printing it.



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 report or work with existing reports through the Reports pane toolbar icons. For more information, see [Reports](#).

Reports pane toolbar



Add: Use this icon to start the process of creating a backup or restore report. 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 icon is the only one that is enabled.



Edit: Click this icon to open the **Report Wizard** and change 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**.

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 is displayed 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.

Backup

- [Backup overview](#)
- [Reviewing backup transport options](#)
- [Determining application consistency](#)
- [Understanding retention policies and space-saving technologies](#)
- [Performing optional configurations](#)
- [Configuring a backup job](#)

Backup overview

vRanger offers two types of backups: virtual and physical. Each type is described in the following topics. Where information relates to only one type of backup, that restriction is noted in the topic title or as a note in the text.

- [Virtual machine \(VM\) backup](#)
- [Physical machine backup](#)

Virtual machine (VM) backup

vRanger protects two varieties of VMs:

- VMware
- Hyper-V

For more information on the different types of virtual backups, see the following topics:

- [Protecting VMware VMs](#)
- [Protecting Hyper-V VMs](#)

Protecting VMware VMs

When protecting a VMware® virtual environment, vRanger uses VMware snapshot technology to store incoming write requests temporarily 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 back up a VMware VM that already has an open snapshot and can back up the open snapshot, but any secondary consolidated helper snapshots are closed prior to running the backup.

With Inventory Node Selection, you can browse the VMware® vCenter™ or VMware® vCloud Director® inventory and select which VMs, groups, or VMware vSphere® vApp(s)™ you want to protect. You can select a VM, folder, resource pool, vApps, VMware® ESXi™ host, data center, or vCenter, and back up all the VMs located under that node in the tree.

VMware Virtual Volumes and Storage Policies

When you protect a VMware environment, vRanger records data regarding the setup. This data includes whether the VM resides on a DRS-enabled cluster, if there are storage policies involved, and the datastore configuration—network-attached storage (NAS), storage-area network (SAN), VMware Virtual Machine File System (VMFS), or Virtual Volumes (VVols). When you perform a restore or replication, vRanger uses the default settings from the original configuration but provides the option of changing these settings. For example, you can restore a backup from a non-VVol datastore to a VVol datastore, or you can restore a backup from a datastore that is governed by one storage policy to a datastore that is governed by a different policy, or no policy.

i | IMPORTANT: vRanger cannot back up physical raw device map (RDM) partitions. You should see the warning: Incompatible drive detected.

Protecting VMware vApps

VMware vSphere® vApp(s)™ are containers—similar to a resource pool—that help you manage and control resources for the VMs contained within. vApps are used to group the individual components of a multi-tier application properly, allowing for application-level resource control and portability. When protecting vApps, vRanger backs up both the VMs and the vApp metadata to maintain the vApp settings and structure upon restore.

Protecting encrypted VMs

VMware vSphere version 6.5 introduced the ability to encrypt virtual machines, including disks and most VM files. Starting with vRanger 7.8, backup, replication, and restore of encrypted virtual machines is supported. In order to properly protect encrypted virtual machines, the following conditions must be met:

- Virtual machine encryption must be properly configured and deployed in your virtual environment. This includes the setup of the key management server and creating of the VM encryption policy.
- When using vRanger with machine based HotAdd to backup encrypted VMs, the vRanger VM must also be encrypted. When using vRanger with VA-based HotAdd to backup and replicate encrypted VMs, the VA must also be encrypted.
- Backup jobs must not use the SAN transport mode. SAN transport is not supported with encrypted VMs.

Protecting Hyper-V VMs

When protecting a Hyper-V® host, the vRanger Hyper-V Agent is installed on the host. The vRanger Hyper-V Agent works with the VSS writer present on all Hyper-V hosts to back up every VM on the host and send the backups directly to the repository. If the host is on the same network as the repository location, such as a SAN, the client can perform LAN-free backups.

i | IMPORTANT: Backing up a Hyper-V VM 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.

Incremental and differential backups with Resilient Change Tracking

Hyper-V 2016 introduces Resilient Change Tracking (RCT), a native utility that provides change tracking for Hyper-V virtual machines. RCT maintains a record of block changes over time, enabling vRanger to quickly identify blocks changed since the last backup (for differentials, since the last full backup) and eliminating the need to scan the disk to identify changes.

i | NOTE: Resilient Change Tracking is enabled by default on Windows Server 2016 and requires no additional configuration in vRanger. Incremental backup jobs for Windows Server 2016 Hyper-V will use RCT by default.

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 runs 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 process can result in network backups or LAN-free backups. For backups to be LAN-free, the source server and repository must exist on the same SAN.

i | IMPORTANT: Physical backups are only supported on Windows® machines matching the list of supported operating systems. For more information, see the *Quest vRanger Installation/Upgrade Guide*.

Reviewing backup transport options

vRanger offers different data transport options for backup tasks.

i | NOTE: The backup transport method describes only how data is read from the source server, not how the data is sent to the repository.

- [VMware backup transport options](#)
- [Hyper-V backup transport options](#)
- [Physical backup transport options](#)
- [Transport selection order: VM backup](#)

VMware backup transport options

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.

- [HotAdd](#)
- [SAN backups](#)
- [LAN](#)

HotAdd

vRanger includes support for the VMware® HotAdd disk transport functionality, which permits LAN-free backups from a VM.

vRanger can use HotAdd for two backup methods:

- **VA-based HotAdd:** Mounts the source VM's disk to the vRanger virtual appliance (VA) deployed on the source host or cluster. This method lets vRanger have direct access to the VM data through the VMware I/O stack rather than the network.
- **Machine based HotAdd:** If vRanger is installed in a VM, this method mounts the source VM's disk to the vRanger VM. This method lets vRanger have direct access to the VM data through the VMware 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 VM—either the vRanger VA or the vRanger VM—must be able to access the target VM's datastores. In addition, all hosts to which the vRanger VM or VA could be vMotioned using VMware vSphere® vMotion® 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 vRanger for HotAdd

When using HotAdd, make sure to disable Automount on the vRanger machine. This step prevents Windows® on the vRanger VM from assigning a drive letter to the target VMDK.

To disable Automount:

- 1 Click **Start > Run**, and enter **diskpart**.
- 2 Run the **automount disable** command to disable automatic drive letter assignment.
- 3 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 1 MB blocks, the proxy must also be on a datastore with 1 MB blocks. For more information, see the VMware® VDDK 1.2.1 Release Notes at <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 VM. This issue might result in the vRanger volume being removed from the VM along with the source VMDK. For more information, see the VMware VDDK 1.2.1 Release Notes at <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 a VMware® ESXi™ 6.0 host cannot HotAdd a disk from an ESXi 5.1 host.

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 protect must also be properly zoned and 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:
Click **Start > 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 VMware® ESXi™ servers, data is sent by way of the VMware VDDK transport.

i | NOTE: Generally speaking, network backups and restores are slower than LAN-free backups and restores. If your environment supports LAN-free backups, use that option whenever possible.

Transport selection order: VM backup

When backing up VMs 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 VAs 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 following topics describe the steps taken to select an automatic transport.

- [vRanger: installed on a VM](#)
- [vRanger: installed on a physical machine](#)

vRanger: installed on a VM

When installed on a VM, vRanger checks the possible transports in the following order, selecting the first option available to ensure the best performance possible.

- 1 **VA on source host:** vRanger first checks the source host for a VA. 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.

i | IMPORTANT: If a VA is configured and HotAdd is not available, a network backup is performed from the VA. No further transport checks are performed.

- 2 **VA on source cluster:** If the source host does not have a VA, vRanger determines whether the host is a member of a cluster. If so, vRanger checks for a shared VA 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.

i | IMPORTANT: If a VA is configured and HotAdd is not available, a network backup is performed from the VA. No further transport checks are performed.

- 3 **Machine-based HotAdd:** If there is no VA configured, vRanger determines whether 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 VA 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 following order, selecting the first option available. The preferred transports are checked first, ensuring the best performance possible.

- 1 **VA on source host:** vRanger first checks the source host for a VA. 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.

i | IMPORTANT: If a VA is configured and HotAdd is not available, a network backup is performed from the VA. No further transport checks are performed.

- 2 **VA on source cluster:** If the source host does not have a VA, vRanger determines whether the host is a member of a cluster. If so, vRanger checks for a shared VA 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.

i | IMPORTANT: If a VA is configured and HotAdd is not available, a network backup is performed from the VA. No further transport checks are performed.

- 3 **Machine-based SAN:** If there is no VA configured, vRanger determines whether the vRanger server is configured for SAN backups.
- 4 **Machine-based LAN:** If a VA 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[®] VMs, 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 process can result in network backups or LAN-free backups. 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 VM backups and backups of physical servers. For more information, see the following topics:

- [Application consistency for virtual backups](#)
- [Application consistency for physical backups](#)

Application consistency for virtual backups

vRanger provides various levels of protection for VMs. The means that each level uses to accomplish its task depends on the type of VM being protected.

- [Application consistency for VMware VMs](#)
- [Application consistency for Hyper-V VMs](#)

Application consistency for VMware VMs

By default, vRanger does not provide quiescing during backups. When you enable the **Enable Guest Quiescing** option, quiescing in vRanger is provided by leveraging the VMware® Tools installed in the VM.

This feature can provide three different levels of backup 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.
- **File-system consistent:** File-system consistency is achieved through standard quiescing by using 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 prior to creating the VM snapshots. This option 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 depends on the version of VMware® ESXi™—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 1. Achieving consistency

File-level quiescing			
	ESXi 6.0	ESXi 6.5	ESXi 6.7
Windows Server 2008 R2 SP1	VMware VSS	VMware VSS	VMware VSS
Windows Server 2012	VMware VSS	VMware VSS	VMware VSS
Windows Server 2012 R2			
Windows Server 2016	VMware VSS	VMware VSS	VMware VSS
Application-level quiescing			
	ESXi 6.0	ESXi 6.5	ESXi 6.7
Windows Server 2008 R2 SP1	vzShadow.exe	vzShadow.exe	vzShadow.exe
Windows Server 2012	vzShadow.exe	vzShadow.exe	vzShadow.exe
Windows Server 2012 R2			
Windows Server 2016	vzShadow.exe	vzShadow.exe	vzShadow.exe

As shown in the preceding table, application consistency is not always available with the basic quiescing options. In these situations, you may use the vRanger VSS Tools—vzShadow.exe—for application-level consistency.

IMPORTANT: Due to the database architecture, vRanger cannot take application-consistent backups of vCenter Server Appliance (VCSA) VMs. Use caution when backing up a VCA VM to ensure reliable recovery.

IMPORTANT: Application-consistent backups are not available for encrypted VMs.

Application consistency for Hyper-V VMs

vRanger can provide two levels of consistency for Hyper-V® VM backups: file-system consistency 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 the Microsoft Volume Shadow Copy Services (VSS) to put supported applications into a consistent state during a backup. This option 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 VM 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, see [https://msdn.microsoft.com/en-us/library/dd405549\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/dd405549(v=vs.85).aspx).

Application consistency for physical backups

vRanger can provide two levels of consistency for physical backups: crash consistent 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 shutdown.
- **Application consistent:** vRanger engages the Microsoft VSS to put supported applications into a consistent state during a backup. This option 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—through 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 number of savepoints retained 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.

i | 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 several savepoints that are grouped for two reasons: The first is to implement the retention policy; and the second is to maintain the parent-child relationships that exist when you choose to use differential or incremental backups. The backup set for each type of backup is different.

- **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 in retention terms means 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 following scenario 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 backup jobs 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 following scenario depicts a daily incremental backup job, with a Savepoint Count of 7 and a Threshold Count of 6:

Table 2. 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 following scenario depicts a daily differential backup job, with a Savepoint Count of 7, a Threshold Count of 6, and a Threshold Size of 50%:

Table 3. 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.

Table 3. Differential backup retention example

Day	Action
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.
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 set up a backup job, you may want to make the following optional vRanger configurations:

- [Enabling VMware Changed Block Tracking \(CBT\)](#)
- [Adding a custom backup group](#)
- [Configuring VSS tools for VMware VMs](#)


Enabling VMware Changed Block Tracking (CBT)

VMware® 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.

VMware vSphere® supports CBT, and most VMs running in this environment can use it. The VMs must be Hardware Version 7 or later, and have been created and hosted in VMware® ESXi™ 6.0 or later hosts. VMs that are created in VMware® ESX® 4 or earlier must be migrated to Hardware Version 7 or later for CBT to be supported. CBT must be enabled for each VM with which CBT is to be used.

i | **NOTE:** VMs 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.

Change Block Tracking is enabled or disabled on a per-VM basis. There are two methods to enable CBT:

- In the **Options Selection** page of the vRanger Backup Wizard, select **Enable Change Block Tracking**.
i | **NOTE:** This option is enabled by default. If you are upgrading a previous version of vRanger, existing jobs will use their current settings.
- In the **My Inventory** view, right-click the VM for which want to change the CBT setting, and then select **Enable Change Tracking**. The **Changed Block Tracking Icon**  appears next to any VM on which CBT is enabled.

Adding a custom backup group

This feature lets you organize physical servers, Hyper-V® servers, or VMware® vCenter™ objects—VMs, hosts, and folders—into groups, which lets you create a backup job for only the selected nodes. This feature provides additional flexibility for creating backup jobs.

i | 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 In 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**.
 - 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 group.
- 4 Click **Add**.

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 VMs

VSS Tools are only available for VMware® 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 topics describe the procedures for the two available methods of configuring VSS Tools with vRanger.

- [Configuring VSS tools using the automated method](#)
- [Configuring a backup job](#)

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 VMs. This optional method uses a small executable tool—**vzShadow.exe**—deployed to the source VM and is enabled by VMware Tools during quiescing. The vRanger VSS Tools must be deployed to each VM for which you need application-consistent backups. For more information, see [Application consistency for virtual backups](#).

i | NOTE: This method is available only for powered-on VMs running Windows Server® 2008 R2 SP1 or later.

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 When the **Configure VSS Tools** dialog box appears, enter the name and password for an account with administrative privileges on the VM, and click **Next**.

If the selected VM has more than one disk, the available disks are shown in the **Disk for VSS snapshot section**.
- 3 Select the disks to quiesce.
- 4 If preferred, select **Perform application log truncation**.
- 5 Click **Configure**, and click **Next**.

vRanger performs the following actions:

- Creates the **C:\Program Files\Quest\VRanger\VSS** directory, and populate the directory with the **vzShadow** executable and supporting files.
- Creates the **C:\Program Files\VMware\VMware Tools\backupscripsts.d** directory, and creates 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 When the **Configure VSS Tools** dialog box appears, enter the name and password for an account with administrative privileges on the VM, and click **Next**.
- 3 Click **Remove**.

The files are removed from the VM, but the directories remain.

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 topics that follow.

To configure a backup job, complete the following tasks:

- [Starting the Backup Wizard](#)
- [Excluding VMs](#)
- [Including hard disks of physical machines or VMs](#)
- [Selecting a repository](#)
- [Selecting a backup data transport method for VMware machines](#)
- [Configuring backup options](#)
- [Selecting a Retention Policy](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the backup job](#)

Starting the Backup Wizard

The first step to creating a vRanger backup job is to select the inventory node—physical server, VMware® vCenter™, host, backup group, and so on—that you want to back up. After selecting the node, you may start the Backup Wizard as described in the following procedures.

You may also create a custom grouping consisting of multiple physical servers or of existing objects from multiple hosts or vCenters. For more information, see [Adding a custom backup group](#). Backup groups are treated the same as any other inventory node when it comes to creating backup jobs.

i | 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](#)

i | **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](#), starting the Backup Wizard creates a job to back up the selected inventory node or backup group. If a node is not selected, the Add icon on 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 select the object or backup group.

i | **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 selected node and select **Backup <object name>**.
 - With the preferred inventory node selected, click the **Add** drop-down menu, and then click **Backup Job**.

The **Backup Wizard** starts.

- 3 In the **Job Name** field, enter a name for the backup job.
- 4 [Optional] In the **Job Description** field, enter a description for the backup job.
- 5 Click **Next**.

i | **NOTE:** If your selected inventory node contains only one server, skip to [Including hard disks of physical machines or VMs](#). If your selected node contains multiple servers (such as a backup group or host), proceed to [Excluding VMs](#).

Selecting a node in the My Jobs view

You may also start 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**

i | **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**.
- 3 On the **Inventory Node Selection** page, click the **VMware**, **Hyper-V**, **Physical**, or **vCloud** tab.
- 4 Select the node you want to back up.

i | **TIP:** 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** field, enter a name for the backup job.
- 6 [Optional] In the **Job Description** field, enter a description for the backup job.
- 7 Click **Next**.

Excluding VMs

i | **NOTE:** If you selected a node that includes multiple servers—such as a backup group or host—and multiple VMs, the next page in the Backup Wizard is the **Virtual Machine Exclusion** page. This page is 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 appears if you created a job for a node that includes multiple VMs, 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. You then use this page to select the two VMs you do not want to include in the job.

To exclude VMs:

- 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 machines or VMs

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 VM, the Current Virtual Machines pane lists the servers included in the backup job, based on your selections in the **Excluding VMs** 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 preceding example, the six-disk server receives an extra 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 select all potential disks; selecting this option ensures that any disk added to a server in the future is included in the backup job.

To include hard disks of physical or VMs:

- 1 In the **Included Hard Disks** pane, select the disks that vRanger should include in the rule it applies to each server.

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 clear selected disks for the backup group. However, you cannot clear Disk 1 for the group.

- 2 [Optional] 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 repository from this page. For more information, see [Adding a repository](#).

vRanger supports the following repository options for storing backup archives:

i | **NOTE:** Any storage device or appliance that supports CIFS or NFSv3 is expected to be compatible with vRanger.

- **Windows Share (CIFS)**
- **NFS**
- **NetVault SmartDisk (NVSD):** The Quest 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://quest.com/products/netvault-smartdisk/> or the *Quest vRanger Integration Guide for Quest NetVault SmartDisk*.
- **EMC Data Domain Boost (DDB):** Integrating EMC Data Domain Boost (DD 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 the *Quest vRanger Integration Guide for EMC Data Domain Boost (DD Boost)* or <http://www.emc.com/data-protection/data-domain/data-domain-boost.htm>.

- **Quest Rapid Data Access (RDA):** Provided by the Quest DR Series appliances—purpose-built, disk backup appliances that use Quest deduplication technology to improve backup and recovery processes. For more information on Quest DR Series appliances, see <http://quest.com/products/dr-series-disk-backup-appliances/> or the *Quest vRanger Integration Guide for Quest DR Series Disk Backup Appliance*.

i | **NOTE:** Quest RDA and EMC DDB repositories support repository replication. Only one repository can be configured per backup job. For more information on configuring repository replication, see [Managing repository replication](#).

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 repository, see [Adding a repository](#) for more information.

i | **NOTE:** When using Quest RDA or EMC DD Boost repositories, you may configure the repository to replicate savepoints to another like device. For more information, see [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 select the best transport method at run-time automatically; you can also configure the transport manually by using the [Using Custom transport selection](#) option.

i | **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 a backup data transport method for VMware machines, complete one of the following procedures:

- [Using Automatic transport selection](#)
- [Using Custom transport selection](#)

Using 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: VM backup](#).

To use the Automatic Transport selection:

- 1 On the **Transport Selection** page, select **Automatic transport selection**.

The selected transport method appears at the bottom of the page.

i | **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method might not always be what is shown on the page during the job setup.

vRanger verifies whether any VAs 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 \(VA\)](#).

- 2 Click **Next**.

Using 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 option enables the custom settings.

- 2 Select where you want the job to occur from the following options:
 - On this vRanger machine.
 - On a vRanger VA.

vRanger verifies whether any VAs 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 \(VA\)](#).

NOTE: Using the vRanger VA 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**.

CAUTION: 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, select your preferences from the options described in the following table.

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 might not have changed significantly since the last backup, making another backup redundant.	<ul style="list-style-type: none">• 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.	<ul style="list-style-type: none">• VMware• Hyper-V• Physical
Compress backed up files.	This option prompts vRanger to compress the archive. NOTE: The following repositories do not support compression: Quest NetVault SmartDisk, EMC® Data Domain® (DD Boost™), and Quest RDA.	<ul style="list-style-type: none">• VMware• Hyper-V• Physical
Update notes with latest backup results.	If this flag is selected, vRanger updates the VM Notes in VMware® vCenter™ with the most recent backup status. vRanger maintains the existing VM notes in vCenter, appending the backup status to the end.	<ul style="list-style-type: none">• 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. For more information, see Application consistency for virtual backups .	<ul style="list-style-type: none">• VMware

Option	Description	Inventory type
Enable VSS Application-level quiescing	This option engages the Microsoft 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® VMs, selecting Enable guest quiescing enables VSS application quiescing if supported by the Guest OS and application.	<ul style="list-style-type: none"> Physical
Enable VSS log truncation	This option uses the Microsoft 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 VMs, selecting Enable guest quiescing enables VSS log truncation if supported by the Guest OS and application.	<ul style="list-style-type: none"> Physical
Enable Active Block Mapping (ABM)	Selecting this option enables ABM, 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: <ul style="list-style-type: none"> 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, the page files are backed up by default. If you are upgrading to vRanger 7 from version 6.1, page files are backed up unless you edit each job and explicitly select the option to skip page files.	<ul style="list-style-type: none"> 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 .	<ul style="list-style-type: none"> VMware Hyper-V Physical
Enable Change Block Tracking (CBT)	VMware 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. For more information, see Enabling VMware Changed Block Tracking (CBT)	<ul style="list-style-type: none"> VMware

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](#).

i **IMPORTANT:** For backup jobs configured to use a replicated repository, the chosen retention policy is applied to both the source and target repositories. For more information, see [Managing repository replication](#).

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 VM and detects the blocks actively used by the disk, as opposed to blocks that were 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.

i NOTE: In vRanger 6.1, page files and hibernation files were not backed up by default. In vRanger 7, page and hibernation files are backed up by default. If you are upgrading to vRanger 7 from version 6.1, page files are backed up unless you edit each job and explicitly select the option to skip the page file or hibernation file, or both.

ABM does not back up 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 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 VM and detects the blocks actively used by the disk, as opposed to blocks that were 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.

i NOTE: In vRanger 6.1, page files and hibernation files were not backed up by default. In vRanger 7, page and hibernation files are backed up by default. If you are upgrading to vRanger 7 from version 6.1, page files are backed up unless you edit each job and explicitly select the option to skip the page file or hibernation file, or both.

CAUTION: 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):** 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 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® VM 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 setting is 7, the maximum value is 1000.
 - 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 option defines the number of incremental backups to perform before a full backup is done.
Use the up and down arrows to select a threshold count—the default is 6.
 - **Full Backup Days:** 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 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 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 VM and detects the blocks actively used by the disk, as opposed to blocks that were 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, page and hibernation files are backed up by default. If you are upgrading to vRanger 7 from version 6.1, page files are backed up unless you edit each job and explicitly select the option to skip the page file or hibernation file, or both.

CAUTION: 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):** 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 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® VM 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 setting is 7, the maximum value is 1000.
- 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 option defines the number of incremental backups to perform before a full backup is done.
Use the up and down arrows to select a threshold count—the value is 6.
 - **Full Backup Days:** 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 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 might 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 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	<p>In this section, you establish how often the report should be generated. There are five options within this section:</p> <ul style="list-style-type: none"> • Daily: The daily option can be scheduled to run the report every weekday or every <i>x</i> 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: <ul style="list-style-type: none"> ▪ Day [<i>x</i>] of every [<i>y</i>] month: <p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the job occurs.</p> <p><i>y</i> 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.</p> ▪ The [<i>f</i>] [<i>d</i>] of every [<i>y</i>] month(s): <p><i>f</i> can be either: first, second, third, fourth or last.</p> <p><i>d</i> can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p> <p><i>y</i> 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.</p> • Yearly: The yearly option offers the following configurations: <ul style="list-style-type: none"> ▪ Every [<i>m</i>] [<i>x</i>]: <p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p> <p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the report occurs.</p> ▪ The [<i>f</i>] [<i>d</i>] of [<i>m</i>]: <p><i>f</i> can be either: first, second, third, fourth, or last.</p> <p><i>d</i> can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p> <p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p> • Interval: The interval option lets you select the number of days, hours, and minutes that should pass between backup jobs. <p>NOTE: The interval selected must be greater than or equal to five minutes.</p>
Range of Recurrence	<p>Select the starting date and ending date:</p> <ul style="list-style-type: none"> • 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.

i | **TIP:** 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**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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](#)
- [Performing a full restore for VMware VMs](#)
- [Performing a full restore for Hyper-V® VMs](#)
- [Performing a full restore for VMware vApps](#)
- [Performing a full restore of a physical machine](#)
- [Performing an FLR on Windows](#)
- [Performing an FLR on Linux](#)
- [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 topics:

- [Full restore](#)
- [Restoring a file](#)
- [Restoring from a manifest file](#)

If you are restoring a VMware® VM, be aware of the following regarding the destination that you target for a restore:

- You can change the storage-destination settings. Although vRanger uses the defaults from the original configuration, you change the host and storage targets. For example, you can restore a backup from a non-VVol datastore to a VVol datastore. In another example, you can restore a backup from a datastore that is governed by one storage policy to a datastore that is governed by a different policy, or no policy.
- You can select different storage for the target of the primary VM file versus the disks.
- If the target uses DRS, you must select the DRS, which governs which datastore in the cluster is used at restore time; you cannot select a specific datastore.
- If your environment uses storage policies, datastores are listed under the applicable policy. Any datastores not associated with a storage policy are listed separately.

Full restore

Using vRanger, recovery of an entire virtual machine (VM)—or resource pool, VMware vSphere® vApp(s)™, and so on—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 merges 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 vApp, vRanger restores the VM data and the parent object's metadata, allowing you to restore the entire object state and settings and the actual VM data.

During the normal restore process, information is pulled from the vRanger database to complete the restore job. If 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 an encrypted VMware VM

vRanger leverages vSphere Virtual Machine Encryption to protect and restore encrypted VMs. When restoring an encrypted VM, you can restore it as either encrypted or unencrypted.

In addition, vRanger provides the option to encrypt a restore of an unencrypted backup.

i | **IMPORTANT:** When restoring an encrypted VM while using a vCenter account without administrator privileges, you must configure that account with the appropriate permissions. Refer to the section "Configuring vCenter Permissions" in the *vRanger Installation/Upgrade Guide*.

Restoring a physical server

vRanger restores physical savepoints to physical target servers using the vRanger Restore image—by using a CD or USB drive. This restore image provides a temporary boot environment into which the vRanger tools are loaded and run. 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 VM. A full restore returns the server or servers 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 merges 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.

i | **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 might prevent the restored server from functioning.

For the procedures on restoring a physical savepoint, see [Performing a full restore of a physical machine](#).

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 restore (FLR) 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 an FLR using Catalog Search](#).

Restoring from a manifest file

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 that are not part of a repository—those savepoints that have been recovered from tape, for example.

The Restore from Manifest function can be used for full restores of VMs and physical servers, and FLR. To restore a savepoint using the manifest file, see [Restoring from manifest](#).

Performing a full restore for VMware VMs

vRanger performs a full restore of a 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 [Performing a full restore of a physical machine](#).

To perform a full restore for a VM, complete the following tasks:

- [Starting the Restore Wizard](#)
- [Selecting a VM Name & Folder](#)
- [Selecting a network](#)
- [Selecting a restore data transport method for](#)
- [Configuring restore options](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the restore job](#)

Starting the Restore Wizard

When restoring a server from a savepoint, the Restore Wizard guides you through the steps necessary to complete the restore.

i | **NOTE:** Some of the steps of the Restore Wizard differ depending on the number of machines being restored. Where appropriate, steps separated between single and multiple machines.

To use the Restore Wizard:

- 1 In the main vRanger UI, 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, find the machine to restore. You can filter the machine results by inventory type, and/or search for any part of the machine name.
- 4 You can automatically restore the machine to the latest savepoint. To do so, right-click on the machine row and click **Restore Latest**.
To select a different savepoint, expand the machine row and select the desired savepoint. Right-click the savepoint, and click **Restore**, or click **Restore** on the vRanger toolbar. The Restore Wizard starts.
- 5 On the **Job Information** page, enter the **Job Name** and **Job Description**.
- 6 Click **Next**.

Selecting a VM Name & Folder

The VM Name & Folder page lets you specify the name for the restored VM and configure the datacenter to which the savepoints are to be restored.

To select a VM and datacenter:

- 1 For a single VM:

- a In the **Virtual Machine Name** field, accept the existing VM name, or enter a new name.
 - b Select the appropriate datacenter or VM folder from the available options.
 - c Click **Next**.
- 2 For multiple VMs:
 - a For each machine, the **Restore As** field displays the name that will be given to the restored VM. To restore the VM with the original name, do nothing.
OR
To restore the VM with a different name, click on the name in the Restore As column, and edit the VM name as desired.
 - b For each machine, confirm the datacenter or VM folder listed in the VM Folder field.
OR
Click the browse button and select a datacenter or VM folder from the available options. Click **OK**.
 - c Click **Next**.

Selecting a Compute Resource

The **Compute Resource** page allows you to specify the computing resource (cluster, host, vApp, or resource pool) that will run the restored virtual machine.

To select a computing resource:

- 1 For a single VM:
 - a In the right-hand pane, select the appropriate compute resource.
 - b Click **Next**.
- 2 For multiple VMs:
 - a For each machine, confirm the resource listed in the **Compute Resource** field.
OR
Click the browse button and select a **Compute Resource** from the available options. Click **OK**.
 - b Click **Next**.

Selecting Storage

The Storage page allows you to configure the primary datastore which acts as the default selection for restored VM disks. You may also deselect individual disks, or map individual disks to a storage device other than the primary datastore.

To configure storage selections:

- 1 For a single VM:
 - a In the **Primary Datastore** field, confirm the datastore listed, or click the browse icon and select a different datastore from the list.

This page lists all datastores associated with the selected host. Datastores are first grouped based on whether they are governed by a storage policy. Datastores are also grouped based on whether they reside in a DRS cluster. You can expand or collapse groups as needed.

By default, vRanger assigns the same storage selection to the primary VM and its disks, if applicable.

- b To select a different datastore for a disk, click the browse button in the **Target Datastore** column entry for the disk and select the applicable datastore or DRS. Click **OK**.

i **NOTE:** VVol datastores only support thin disks. If your VM resided on a thick disk during the backup and you select a VVol datastore for the restored VM, vRanger automatically converts the setup to thin disks.

- c Click **Next**.

2 For multiple VMs:

- a For each machine, confirm the datastore listed in the **Primary Storage** field.
OR
Click the browse button and select the appropriate storage from the available options. Click **OK**.
- b Click **Next**.

Selecting a network

The **Network Selection** page lets you configure the virtual network to which each vNIC should be connected upon restore.

To select a network:

- 1 For a single VM:
 - a For each network adaptor, select the appropriate network from the drop-down list in the **Restore to Network** column.
 - b Click **Next**.
- 2 For multiple VMs:
 - a Expand the VM entry to view the network adaptors that are part of that VM.
 - b For each network adaptor, select the appropriate network from the drop-down list in the right-most column.
 - c Click **Next**.

Selecting a restore data transport method for

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 select the best transport method at run-time; you can also configure the transport manually by using the [Using Custom transport selection](#) option.

- i** **NOTE:** The Transport Selection page is not available when restoring a physical or Hyper-V® virtual server. When restoring physical servers and Hyper-V VMs, vRanger sends the data to restore over the network from the repository. Because this method is the only one 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 Automatic transport selection](#)
- [Using Custom transport selection](#)

Using 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: VM backup](#).

To use Automatic transport selection:

- 1 On the **Transport Selection** page, select **Automatic transport selection**.

The selected transport method appears at the bottom of the page.

i | IMPORTANT: Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method might not always be what is shown on the page during the job setup.

vRanger verifies whether any VAs 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 \(VA\)](#).

- 2 Click **Next**.

Using 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 Custom transport selection:

- 1 On the **Transport Selection** page, select **Custom transport selection**.

This option enables the custom settings.

- 2 Select where you want the job to occur from the following options:
 - On this vRanger machine.
 - On a vRanger VA.

vRanger verifies whether any VAs 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 \(VA\)](#).

i | NOTE: Using the vRanger VA 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**.

i | CAUTION: 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 option overwrites any existing disk or config files of the same name as those files in the destination host.
 - **Force power on:** This option 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.

i | **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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking **Finish**, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 full restore for Hyper-V[®] VMs

vRanger performs a full restore of a 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 [Performing a full restore of a physical machine](#).

To perform a full restore for a VM, complete the following tasks:

- [Starting the Restore Wizard](#)
- [Selecting a VM Name & Folder](#)
- [Selecting a network](#)
- [Selecting a restore data transport method for](#)

- [Configuring restore options](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the restore job](#)

Starting the Restore Wizard

When restoring a server from a savepoint, the Restore Wizard guides you through the steps necessary to complete the restore.

i | **NOTE:** Some of the steps of the Restore Wizard differ depending on the number of machines being restored. Where appropriate, steps are separated between single and multiple machines.

To use the Restore Wizard:

- 1 In the main vRanger UI, 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, find the machine to restore. You can filter the machine results by inventory type, and/or search for any part of the machine name.
- 4 You can automatically restore the machine to the latest savepoint. To do so, right-click on the machine row and click **Restore Latest**.

To select a different savepoint, expand the machine row and select the desired savepoint. Right-click the savepoint, and click **Restore**, or click **Restore** on the vRanger toolbar. The Restore Wizard starts.
- 5 On the **Job Information** page, enter the **Job Name** and **Job Description**.
- 6 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 apply to all VMs, do the following:
 - a Select the host from the **Master Host** drop-down list.
 - b Confirm that the path that appears in the **Virtual Machine Location** field is correct. To change this path, use the browse (...) button to locate and select a new location.
 - c Click the **Master Disk Storage Datastore** browse button to locate and select the datastore to use for all VMs. The Master Datastore lets you select all datastores simultaneously.
 - d Click **Assign**.
- 2 To select options for individual Hyper-V® VMs, do the following:
 - a Select the disks you want to restore. Disks disabled for Restore appear in italics.
 - b 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.
 - c To select a different datastore for each VM, 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 network adapters should be connected upon restore. You can configure a Master Network, which will be pre-populated as the default selection for each network adaptor in the restore job.

To select a network:


- 1 Do one of the following:
 - To use the same network for all VMs, select a host from the **Master Network** drop-down list to which all VMs should be restored. Click **Assign**.
 - 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.
- 2 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 option overwrites any existing disk or config files of the same name as those files in the destination host.
 - **Force power on:** This option 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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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**.

Performing a full restore for VMware vApps

vRanger performs a full restore of a VMware vSphere® vApp(s)™ 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 [Performing a full restore of a physical machine](#).

To perform a full restore for a vApp, complete the following tasks:

- [Starting the Restore Wizard](#)
- [Selecting a VM Name & Folder](#)
- [Mapping storage](#)
- [Selecting a restore data transport method for](#)
- [Configuring restore options](#)

- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the restore job](#)

Starting 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 UI, 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 You can automatically restore the machine to the latest savepoint. To do so, right-click on the machine row and click **Restore Latest**.
To select a different savepoint, expand the machine row and select the desired savepoint. Right-click the savepoint, and click **Restore**, or click **Restore** on the vRanger toolbar. The Restore Wizard starts.
- 5 On the **Job Information** page, enter the **Job Name** and **Job Description**.
- 6 Click **Next**.

Selecting a destination

The Destination Selection page lets you configure the virtual data center (vDC) to which the VMware vSphere® vApp(s)™ savepoints are to be restored, and 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 list, select the virtual data center to which the vApp should be restored.
- 3 For each network listed, select a network for the restored vApp in the **Restore to Network** column.
- 4 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 list, select the applicable storage profile.
Selecting a storage profile allows VMware® vCloud Director® to select the appropriate datastore automatically based on your storage profile settings.
- 2 To specify a datastore manually, select **Any** in the **Storage Profile** drop-down list, and select the applicable datastore in the **Datastore** drop-down list.
- 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 VMware vSphere® vApp(s)™, the Restore Wizard uses the **Automatic transport selection** option by default to select the best transport method automatically at run-time. You can also configure the transport manually by using the [Using Custom transport selection](#) option.

i **NOTE:** The Transport Selection page is not available when restoring a physical or Hyper-V® virtual server. When restoring physical servers and Hyper-V VMs, vRanger sends the data to restore over the network from the repository. Because this method is the only one 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 Automatic transport selection](#)
- [Using Custom transport selection](#)

Using 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: VM backup](#).

To use Automatic transport selection:

- 1 On the **Transport Selection** page, select **Automatic transport selection**.

The selected transport method appears at the bottom of the page.

i **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method might not always be what is shown on the page during the job setup.

- 2 Click **Next**.

Using 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 Custom transport selection:

- 1 On the **Transport Selection** page, select **Custom transport selection**.

This option enables the custom settings.

- 2 Select where you want the job to occur from the following options:
 - On this vRanger machine.
 - On a vRanger VA.

vRanger verifies whether any VAs 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 \(VA\)](#).

NOTE: Using the vRanger VA 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**.

CAUTION: This option causes reduced performance.

5 Click **Next**.

Configuring restore options

The Options Selection page displays the restore options available for VMware vSphere® vApp(s)™ restore jobs.

To configure restore options:

1 On the **Options Selection** page, select from the following restore options:

- **Force overwrite:** This option overwrites any existing disk or config files of the same name as those files in the destination host.
- **Force power on:** This option 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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 full restore of a physical machine

Like the restore process for a VM, the restore process for a physical machine uses the vRanger Job Wizard. While the process of restoring physical machines is similar to the VM 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:

- [Creating the boot media](#)
- [Booting the target server](#)
- [Selecting a savepoint](#)
- [Mapping disks](#)
- [Mapping disks](#)
- [Setting the recurrence schedule](#)

- [Configuring email notifications](#)
- [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. You can download the restore image files—**vzRestore-*<version>*.iso** and **vzRestore-*<version>*.img**—from the Quest website.

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 following procedures 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 Download the **vRanger BMR Restore Image (ISO CSD)** file from <http://support.quest.com/vranger/>
- 3 If necessary, copy the **vzRestore-*<version>*.iso** file to the machine with the writable CD/DVD drive.
- 4 Right-click the ISO file, select **Open with**, and then click **Windows Disk Image Burner**.
- 5 From the disk burner drop-down list, select the drive containing the blank disk.
- 6 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 following procedure 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 Download the **vRanger BMR Restore Image (IMG Bootable USB Drive)** file from <http://support.quest.com/vranger/>
- 3 If necessary, copy the **vzRestore-*<version>*.img** file to the machine with the blank USB drive.
- 4 In the Image Writer for Windows UI, browse to the **vzRestore-*<version>*.img** file.
- 5 In the **Device** drop-down list, ensure that the correct USB drive is selected.
- 6 Click **Write**.
- 7 Click **Exit**.

Booting the target server

i | IMPORTANT: The vRanger restore CD is not compatible with UEFI. Please ensure that the target server is not configured to boot into UEFI.

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.

When the boot is complete, a list of available network devices and their connection status appears.
- 2 Do one of the following:
 - If a DHCP server is available, the restore environment obtains its IP configuration automatically, and a prompt appears with the IP information. Record this information to use during the restore job creation process.
 - If no DHCP server is detected, each Ethernet device shows as “Disconnected.” Select the preferred network device. If only one device is present, it is automatically selected. 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**.
- 3 [Optional] If the vRestore environment does not contain the correct drivers for your hardware, enter a Linux® command prompt by pressing **F1**, and add them manually.

i | 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.

When the boot is complete, a list of available network devices and their connection status appears. If no DHCP server is detected, each Ethernet device shows as “Disconnected.”
- 2 Select the preferred network device.

If only one device is present, it is automatically selected.
- 3 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**.

- 5 In the **Setting Static IP** dialog box, 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, enter a Linux[®] command prompt by pressing **F1**, and add them manually.

i | **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 UI, 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.

i | **NOTE:** For the procedures on restoring from a manifest file, see [Performing a full restore from manifest for VMware VMs](#).

- 4 You can automatically restore the machine to the latest savepoint. To do so, right-click on the machine row and click **Restore Latest**.

To select a different savepoint, expand the machine row and select the desired savepoint. Right-click the savepoint, and click **Restore**, or click **Restore** on the vRanger toolbar. The Restore Wizard starts.

- 5 On the Job Information page, enter the Job Name and Job Description.
Click **Next**.

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** field, enter the IP for the target server; for more information, 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 appears.

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	<p>In this section, you establish how often the report should be generated. There are five options within this section:</p> <ul style="list-style-type: none">• Daily: The daily option can be scheduled to run the report every weekday or every <i>x</i> 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:<ul style="list-style-type: none">▪ Day [<i>x</i>] of every [<i>y</i>] month:<p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the job occurs.</p><p><i>y</i> 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.</p>▪ The [<i>f</i>] [<i>d</i>] of every [<i>y</i>] month(s):<p><i>f</i> can be either: first, second, third, fourth or last.</p><p><i>d</i> can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p><p><i>y</i> 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.</p>• Yearly: The yearly option offers the following configurations:<ul style="list-style-type: none">▪ Every [<i>m</i>] [<i>x</i>]:<p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p><p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the report occurs.</p>▪ The [<i>f</i>] [<i>d</i>] of [<i>m</i>]:<p><i>f</i> can be either: first, second, third, fourth, or last.</p><p><i>d</i> can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p><p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p>• Interval: The interval option lets you select the number of days, hours, and minutes that should pass between jobs. <p>NOTE: The interval selected must be greater than or equal to five minutes.</p>
Range of Recurrence	<p>Select the starting date and ending date:</p> <ul style="list-style-type: none">• 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 an FLR on Windows

You can restore a file from a savepoint by accessing the **File Level Restore** 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 select the savepoint and then click the FLR icon on 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 FLR 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.

- i** | **NOTE:** If the volume was created on Windows Server® 2012 or later, the vRanger machine must also use Windows Server 2012 or later to complete FLR. Older systems may not show data on GPT disks.
- VMs with dynamic disks are not supported for FLR.

Performing an FLR 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 box.

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.

- i** | **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 FLR using Catalog Search:

- 1 On the main vRanger UI, navigate to the **My Repositories** pane, and then select the repository in which you want to search.
- 2 Do one of the following to start 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 box appears.

Entering a search string in the Catalog Search Criteria field searches all repositories and savepoints for the string.

- 3 To refine the scope of the search further, click **Advanced**.

The **Advanced** option lets you limit the search to a repository or VM (savepoint).
- 4 Select the repository or VM you want to search, enter the search string, and click **Search**.
- 5 Select the preferred savepoint, and click **FLR** for File Level Restore, or **Restore** for the Full Restore.

- i** | **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; a short search string—for instance, “dot”—could result in very slow searching, and SQL Server Express could run out of memory. To minimize performance issues during Catalog Search, make the search string as specific as possible.

Performing a manual FLR

The dialog box for this method of restoring files is made up of the following three panes:

- **Savepoint pane:** In this area, all 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 box contains data after the restore process starts. It includes a column that indicates the percentage of the job that has completed. The dialog box 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 appears.

To perform a manual FLR:

- 1 In the main vRanger UI, 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.

CAUTION: 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.
- 8 Make sure that the status is Finished and Percent Finished is 100%.
- 9 To view the restored file or folder, click the **View in Explorer** link.

Performing an FLR on Linux

FLR from Linux[®] server savepoints requires that a vRanger virtual appliance (VA) is deployed and configured for Linux FLR. For instructions on deploying and configuring the vRanger VA, see the *Quest vRanger Installation/Upgrade Guide*.

After you configure the vRanger VA, 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 FLR](#).

Linux FLR limitations

There are several limitations and requirements that apply to FLRs from Linux[®] VMs.

NOTE: The following summarizes the limitations; for more information, see the *Quest vRanger Installation/Upgrade Guide*.

- **Linux:** FLR from Linux servers requires the use of the vRanger VA. For more information, see [Performing an FLR on Linux](#).
- **Linux:** Not all Linux file systems are supported for FLR. For a list of supported Linux platforms and file systems, see the *Quest vRanger Installation/Upgrade Guide*.
- **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 need to show the operating system (OS) type as Linux. If this setting 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 feature 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:

- [Performing a full restore from manifest for VMware VMs](#)
- [Performing a full restore from manifest for a physical machine](#)
- [Performing FLR from manifest](#)

Performing a full restore from manifest for VMware VMs

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.

To perform a full restore from manifest for VMs, complete the following tasks:

i | **NOTE:** The procedures in this topic apply to restoring from a VMware VM savepoint manifest. For information on restoring a Hyper-V VM from manifest, see [Performing a full restore from manifest for Hyper-V VMs](#). For information on restoring a physical server from a savepoint manifest, see [Performing a full restore from manifest for a physical machine](#).

- [Creating a restore job](#)
- [Selecting a savepoint from the manifest](#)
- [Selecting a VM Name & Folder](#)
- [Selecting a Compute Resource](#)
- [Selecting a network](#)
- [Selecting a transport method](#)
- [Configuring restore options](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the restore job](#)

Creating a restore job

Creating a restore job starts the Restore Wizard, which guides you through the steps of completing the job.

To create a restore job:

- 1 In the vRanger UI, click **My Repositories**.
- 2 On the toolbar, click **Restore from Manifest**, and then click **Restore from VMware Savepoint Manifest File**. The Restore Wizard starts.

- 3 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 [Step 1](#) through [Step 4](#) as needed for multiple savepoints.
- 6 Click **Next**.
- 7 On the **Job Information** page, enter the **Job Name** and **Job Description**.
Click **Next**.

Selecting a VM Name & Folder

The VM Name & Folder page lets you specify the name for the restored VM and configure the datacenter to which the savepoints are to be restored.

To select a VM and datacenter:

- 1 For a single VM:
 - a In the **Virtual Machine Name** field, accept the existing VM name, or enter a new name.
 - b Select the appropriate datacenter or VM folder from the available options.
 - c Click **Next**.
- 2 For multiple VMs:
 - a For each machine, the **Restore As** field displays the name that will be given to the restored VM. To restore the VM with the original name, do nothing.

OR

To restore the VM with a different name, click on the name in the Restore As column, and edit the VM name as desired.
 - b For each machine, confirm the datacenter or VM folder listed in the VM Folder field.

OR

Click the browse button and select a datacenter or VM folder from the available options. Click **OK**.
 - c Click **Next**.

Selecting a Compute Resource

The **Compute Resource** page allows you to specify the computing resource (cluster, host, vApp, or resource pool) that will run the restored virtual machine.

To select a computing resource:

- 1 For a single VM:
 - a In the right-hand pane, select the appropriate compute resource.
 - b Click **Next**.
- 2 For multiple VMs:
 - a For each machine, confirm the resource listed in the **Compute Resource** field.

OR

Click the browse button and select a **Compute Resource** from the available options. Click **OK**.

 - b Click **Next**.

Selecting Storage

The Storage page allows you to configure the primary datastore which acts as the default selection for restored VM disks. You may also deselect individual disks, or map individual disks to a storage device other than the primary datastore.

To configure storage selections:

- 1 For a single VM:
 - a In the **Primary Datastore** field, confirm the datastore listed, or click the browse icon and select a different datastore from the list.

This page lists all datastores associated with the selected host. Datastores are first grouped based on whether they are governed by a storage policy. Datastores are also grouped based on whether they reside in a DRS cluster. You can expand or collapse groups as needed.

By default, vRanger assigns the same storage selection to the primary VM and its disks, if applicable.
 - b To select a different datastore for a disk, click the browse button in the **Target Datastore** column entry for the disk and select the applicable datastore or DRS. Click **OK**.

NOTE: VVol datastores only support thin disks. If your VM resided on a thick disk during the backup and you select a VVol datastore for the restored VM, vRanger automatically converts the setup to thin disks.
 - c Click **Next**.
- 2 For multiple VMs:
 - a For each machine, confirm the datastore listed in the **Primary Storage** field.

OR

Click the browse button and select the appropriate storage from the available options. Click **OK**.

 - b Click **Next**.

Selecting a network

The **Network Selection** page lets you configure the virtual network to which each vNIC should be connected upon restore.

To select a network:

- 1 For a single VM:
 - a For each network adaptor, select the appropriate network from the drop-down list in the **Restore to Network** column.

- b Click **Next**.
- 2 For multiple VMs:
 - a Expand the VM entry to view the network adaptors that are part of that VM.
 - b For each network adaptor, select the appropriate network from the drop-down list in the right-most column.
 - c Click **Next**.

Selecting a transport method

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 select the best transport method at run-time; you can also configure the transport manually by using the [Using Custom transport selection](#) option.

i **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 VMs, vRanger sends the backup data over the network to the repository. Because this method is the only one 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 Automatic transport selection](#)
- [Using Custom transport selection](#)

Using 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: VM backup](#).

To use Automatic transport selection:

- 1 On the **Transport Selection** page, select **Automatic transport selection**.

The selected transport method appears at the bottom of the page.

i **IMPORTANT:** Each time the job runs, vRanger detects your configuration and selects the best transport method available at run-time. That method might not always be what is shown on the page during the job setup.

vRanger verifies whether any VAs 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 \(VA\)](#).

- 2 Click **Next**.

Using 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 Custom transport selection:

- 1 On the **Transport Selection** page, select **Custom transport selection**.

This option enables the custom settings.

- 2 Select where you want the job to occur from the following options:

- On this vRanger machine.
- On a vRanger VA.

vRanger verifies whether any VAs 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 \(VA\)](#).

i | **NOTE:** Using the vRanger VA 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**.

i | **CAUTION:** 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 option overwrites any existing disk or config files of the same name as those files in the destination host.
- **Force power on:** This option 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.

i | **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.

Table 1.

Start Time	The time of day at which to run the job.
Recurrence Pattern	<p>In this section, you establish how often the report should be generated. There are five options within this section:</p> <ul style="list-style-type: none">• Daily: The daily option can be scheduled to run the report every weekday or every <i>x</i> 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:<ul style="list-style-type: none">▪ Day [<i>x</i>] of every [<i>y</i>] month:<p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the job occurs.</p><p><i>y</i> 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.</p>▪ The [<i>f</i>] [<i>d</i>] of every [<i>y</i>] month(s):<p><i>f</i> can be either: first, second, third, fourth or last.</p><p><i>d</i> can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p><p><i>y</i> 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.</p>• Yearly: The yearly option offers the following configurations:<ul style="list-style-type: none">▪ Every [<i>m</i>] [<i>x</i>]:<p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p><p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the report occurs.</p>▪ The [<i>f</i>] [<i>d</i>] of [<i>m</i>]:<p><i>f</i> can be either: first, second, third, fourth, or last.</p><p><i>d</i> can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p><p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p>• Interval: The interval option lets you select the number of days, hours, and minutes that should pass between jobs. <p>NOTE: The interval selected must be greater than or equal to five minutes.</p>
Range of Recurrence	<p>Select the starting date and ending date:</p> <ul style="list-style-type: none">• 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 full restore from manifest for Hyper-V VMs

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.

To perform a full restore from manifest for VMs, complete the following tasks:

- i** | **NOTE:** The procedures in this topic apply to restoring from a Hyper-V VM savepoint manifest. For information on restoring a VMware VM from manifest, see [Performing a full restore from manifest for VMware VMs](#). For information on restoring a physical server from a savepoint manifest, see [Performing a full restore from manifest for a physical machine](#).

- [Creating a restore job](#)

- [Selecting a savepoint from the manifest](#)
- [Selecting a destination](#)
- [Selecting a network](#)
- [Configuring restore options](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [Starting the restore job](#)

Creating a restore job

Creating a restore job starts the Restore Wizard, which guides you through the steps of completing the job.

To create a restore job:

- 1 In the vRanger UI, click **My Repositories**.
- 2 On the toolbar, click **Restore from Manifest**, and then click **Restore from Hyper-V Savepoint Manifest File**. The Restore Wizard starts.
- 3 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 [Step 1](#) through [Step 4](#) as needed for multiple savepoints.
- 6 Click **Next**.
- 7 On the **Job Information** page, enter the **Job Name** and **Job Description**.
- 8 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 apply to all VMs, do the following:
 - a Select the host from the **Master Host** drop-down list.

- b Confirm that the path that appears in the **Virtual Machine Location** field is correct. To change this path, use the browse (...) button to locate and select a new location.
 - c Click the **Master Disk Storage Datastore** browse button to locate and select the datastore to use for all VMs. The Master Datastore lets you select all datastores simultaneously.
 - d Click **Assign**.
- 2 To select options for individual Hyper-V® VMs, do the following:
 - a Select the disks you want to restore. Disks disabled for Restore appear in italics.
 - b 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.
 - c To select a different datastore for each VM, 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 network adapters should be connected upon restore. You can configure a Master Network, which will be pre-populated as the default selection for each network adaptor in the restore job.

To select a network:


- 1 Do one of the following:
 - To use the same network for all VMs, select a host from the **Master Network** drop-down list to which all VMs should be restored. Click **Assign**.
 - 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.
- 2 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 option overwrites any existing disk or config files of the same name as those files in the destination host.
 - **Force power on:** This option turns on the VM after the restore.
 - **Create cluster role for restored virtual machine:** This option 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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value 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.

i | **TIP:** 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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 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 feature 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.

i | **NOTE:** The procedures in this topic apply to restoring from a physical machine savepoint manifest. For information on restoring a VM from a savepoint manifest, see [Performing a full restore from manifest for VMware VMs](#).

To perform a full restore from a manifest for a physical machine, complete the following tasks:

- [Creating the boot media](#)
- [Booting the target server](#)
- [Creating a restore job](#)

- Selecting a savepoint from the manifest
- Mapping disks
- Setting the recurrence schedule
- Configuring email notifications
- 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. You can download the restore image files—**vzRestore-*<version>*.iso** and **vzRestore-*<version>*.img**—from the Quest website.

This task is a procedure in [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.

i | **IMPORTANT:** The procedures for creating a bootable CD differ widely depending on the operating system and software used. The following procedures 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 Download the **vzRestore-*<version>*.img** file from <http://quest.com/products/vranger/>.
- 3 If necessary, copy the **vzRestore-*<version>*.iso** file to the machine with the writable CD/DVD drive.
- 4 Right-click the ISO file, select **Open with**, and then click **Windows Disk Image Burner**.
- 5 From the disk burner drop-down list, select the drive containing the blank disk.
- 6 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 following procedure 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 Download the **vzRestore-*<version>*.img** file from <http://quest.com/products/vranger/>.
- 3 If necessary, copy the **vzRestore-*<version>*.img** file to the machine with the blank USB drive.
- 4 In the Image Writer for Windows UI, browse to the **vzRestore-*<version>*.img** file.
- 5 In the **Device** drop-down list, ensure that the correct USB drive is selected.

- 6 Click **Write**.
- 7 Click **Exit**.

Booting the target server

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 [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.

When the boot is complete, a list of available network devices and their connection status appears.

- 2 Do one of the following:
 - If a DHCP server is available, the restore environment obtains its IP configuration automatically, and a prompt appears with the IP information. Record this information to use during the restore job creation process.
 - If no DHCP server is detected, each Ethernet device shows as “Disconnected.” Select the preferred network device. If only one device is present, it is automatically selected. 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**.
- 3 [Optional] If the vRestore environment does not contain the correct drivers for your hardware, enter a Linux® command prompt by pressing **F1**, and add them manually.

i | 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.

When the boot is complete, a list of available network devices and their connection status appears. If no DHCP server is detected, each Ethernet device shows as “Disconnected.”

- 2 Select the preferred network device.

If only one device is present, it is automatically selected.
- 3 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**.
- 5 In the **Setting Static IP** dialog box, 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, enter a Linux[®] command prompt by pressing **F1**, and add them manually.

i | IMPORTANT: Do not reboot the target server or the IP configuration is lost.

Creating a restore job

Creating a restore job starts the Restore Wizard, which guides you through the steps of completing the job.

This task is a procedure in [Performing a full restore from manifest for a physical machine](#).

To create a restore job:

- 1 In the main vRanger UI, click **My Repositories**.
- 2 On the toolbar, click **Restore from Manifest**, and then click **Restore from Physical Savepoint Manifest File**.

The Restore Wizard starts.

- 3 On the **Mode Selection** page, populate the **Job Name** and **Job Description** fields.
- 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 [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 [Step 1](#) through [Step 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 [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** box, enter the IP for the target server; for more information, 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 appears.

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 [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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value 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 [Performing a full restore from manifest for a physical machine](#).

To configure email notifications:

- 1 On the **Email Notification Selection** page, in the **Selected** column, select the address to which notifications should be sent.

i | TIP: 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 [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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | 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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:


i | NOTE: If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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 FLR from manifest

An FLR from manifest lets you move the files that make up a backup to your preferred location restore a VM whose backup files are not in a repository. This process lets you move the files that make up a backup to wherever you want and still perform a restore.

To perform an FLR from manifest:

- 1 Click **My Repositories**.
- 2 On the My Repositories toolbar, click the **FLR from Manifest icon** .
- 3 In the **File Level Restore from Manifest** dialog box, 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.

CAUTION: 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.
- 10 Verify that the status is Finished and Percent Finished is 100%.
- 11 To view the restored file or folder, click **View in Explorer**.

Replicate VMs

- [Understanding replication of VMs](#)
- [Replication modes](#)
- [Creating a replication job](#)
- [Using replication failover](#)

Understanding replication of VMs

vRanger includes integrated replication based on the proven technology of vReplicator, enabling replication of VMware® virtual machines (VMs) both on site and at remote locations for flexible and efficient disaster recovery preparedness. Combine backup and replication jobs to ensure that you meet your organization's recovery time and recovery point objectives. Perform fast VM failover and recovery at all your sites, no matter where they are located.

i **NOTE:** Replication is not available for Hyper-V® VMs.

If you are looking for information on replicating Quest RDA and EMC DD Boost repositories, see [Managing repository replication](#).

A VM is made up of a set of files. 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.

Be aware of the following regarding the destination that you target for a replication:

- You can change the storage-destination settings. Although vRanger uses the defaults from the original configuration, you change the host and storage targets. For example, you can replicate from a non-VVol datastore to a VVol datastore. In another example, you can replicate from a datastore that is governed by one storage policy to a datastore that is governed by a different policy, or no policy.
- You can select different storage for the target of the primary VM file versus the disks.
- If the target uses DRS, you must select the DRS, which governs which datastore in the cluster is used at replication time; you cannot select a specific datastore.
- If your environment uses storage policies, datastores are listed under the applicable policy. Any datastores not associated with a storage policy are listed separately.

The following table lists the set of files replicated by vReplicator:

Table 1. 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.

Table 1. Files replicated

File extension	Description
-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. The following table lists these files and their purposes:

Table 2. 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 if there is a replication failure. This file created on the target VA at the start of the replication process and removed at the end.
.vzundo	This file is a temporary file that records original data of changed blocks since the replication started. One for each VMDK replicated. If there is a replication failure such as a network failure, the vzundo-script can be run to restore files to their original state. These files are created on the VA 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 start of the replication process and removed when disk replication is completed.
.vmdk-abfg.vztemp	Change block filter file. One for each hard disk when Changed Block Tracking (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 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 and 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 (VA)

vRanger supports VMware® ESXi™ replication by way of the vRanger VA, which leverages the VMware® HotAdd disk transport mechanism. After the VAs are configured and deployed, the use of the VA is automatic and transparent. The following lists some key points about replicating with the VA:

- For instructions on deploying and configuring the VAs, see the *Quest vRanger Installation/Upgrade Guide*.
- At a minimum, deploy one VA on every ESXi server that you plan to use as a replication source or target. For hosts in a cluster, you can deploy the VA to only one host in the cluster; the VA is shared among the cluster's other hosts.

VAs must be used in pairs. For example, if you are replicating to a target host or cluster that is using a VA, use a VA on the source host or cluster as well.

- When replicating encrypted virtual machines, the source-side vRanger virtual appliance must also be encrypted. To ensure a successful failover, both the source and target VA should be encrypted. In addition, the VA home folder must be encrypted.
- 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. For more information, see [Replication modes](#).

Communication between the VAs occurs through an SSH tunnel using AES-256 encryption. For more information, see the **Encryption** description in [Major feature list](#).

Additional replication requirements


The following limitations and requirements apply to replication:

- vRanger replication does not operate with servers that are behind a network address translation (NAT) firewall. To replicate through a NAT firewall properly, you must have an IP tunnel in place between two NATed subnets. Contact your ISP provider to see if this option is available to you.
- The VM 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 VMware® 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 and compatibility, see the VMware documentation at <https://www.vmware.com/support/pubs/>.
- 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 does work with links having average packet loss of less than 2%. Replication is not designed to work in replication environments where packet loss exceeds commercially accepted limits.
- Networks having 99% uptime and availability generally provide for good Replication performance.

Replication with user snapshots

When replicating a VM that contains user snapshots, vRanger replicates all 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 following image, 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 VA.

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 issue is due to a limitation in how snapshots are processed by HotAdd.

If you make hardware changes on the source VM, you need to configure the target VM in the same way *before* the next replication pass. You might 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 \(CBT\)](#) or standard [Differential replication](#). Each of these replication modes has the option of 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. For more information, see the following topics:

- [Replication with Changed Block Tracking \(CBT\)](#)
- [Differential replication](#)

Replication with Changed Block Tracking (CBT)

Replication with 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 VMware® ESXi™ version 6.0 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 on an independent disk that is not affected by snapshots.

i | **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 on a VMFS volume, backed by a SAN, iSCSI, or local disk.

i | **NOTE:** Raw device map (RDM) partitions are not considered VMFS volumes.

- The VM must have zero snapshots when CBT is enabled for a clean start.

i | **NOTE:** CBT must be enabled for each VM that you want to replicate. VMware 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, see [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 CBT if at all possible.

Configuring transport failover options

Replication transport actions occur 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. For VMware® ESXi™ hosts, configure the transport order to **VA-based HotAdd > VA-based Network**.

Pre-seeding replication jobs

vRanger 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:

- [Backing up the source VM](#)
- [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 disaster recovery (DR) site, and create a Full backup job.
- 3 Back up the VM locally.
- 4 Copy the backup savepoint, and move it to the DR site.

You can complete this step 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, which contains a .var file and two .metadata files, for the savepoint in question.

- 5 Using the Restore from Manifest function documented in [Performing a full restore from manifest for VMware VMs](#), restore the backup to the target VMware® ESXi™ 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 Start vRanger, and expand the host where the source VM is located.
- 2 Right-click the VM, 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.

i | **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 In the **Select Target Host** dialog box, next to **Preseed VM UUID**, click the browse button, and select the pre-seeded VM.
- 6 When configuring the primary datastore, ensure that 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 is explained in more detail in the topics that follow.

To create a replication job, complete the following tasks:

- [Selecting the source VM](#)
- [Selecting a target host](#)
- [Selecting the replication mode](#)
- [Selecting replication network mapping](#)
- [Selecting hard disks](#)
- [Selecting replication options](#)
- [Setting the recurrence schedule](#)
- [Configuring email notifications](#)
- [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 <VMname>**.

i | **NOTE:** If there is no VA provided for the selected VM, the Virtual Appliance Not Found window notifies you that the replication job is going to fail without a VA. For more information, see [Understanding the vRanger virtual appliance \(VA\)](#).

- 3 Enter the **Job Name** and **Job Description**, and 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, and 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. For more information, see [Pre-seeding replication jobs](#). 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] In the **Select Target Host** dialog box, next to **Preseed VM UUID**, click the browse button, and select the pre-seeded VM.

i | **IMPORTANT:** After a replication job has been saved, you cannot change the target host. If you want to replicate to a different host, you may either use VMware vSphere® vMotion® to migrate the VM to the preferred host, or create a 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**.

i | **NOTE:** Hybrid replication is not supported when replicating to or from VMware® ESXi™. This feature is deprecated for vRanger 7.2 and later.

- 2 Click **Next**.

Selecting replication network mapping

On 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 Select which Hard Disk(s) to replicate page lets you include or exclude specific disks from replication.

To select hard disks:

- 1 Select the hard disks you want to replicate.
You may select a **Primary Datastore**, which is used as the default for all disks in this job, letting you configure only the exceptions.
- 2 In the **Target Datastore** column, verify that the Primary Datastore has been populated.
- 3 If you want the disk to replicate to a different datastore, click the browse button (...) next to the **Target Datastore** selection.
- 4 On the **Select Primary Storage** page, review the list of datastores.

This page lists all datastores associated with the selected host. Datastores are first grouped based on whether they are governed by a storage policy. Datastores are also grouped based on whether they reside in a DRS cluster. You can expand or collapse groups as needed.

i **NOTE:** VVol datastores only support thin disks. If your VM resided on a thick disk during the backup and you select a VVol datastore for the replicated VM, vRanger automatically converts the setup to thin disks.

- 5 Select the applicable datastore, and click **OK**.
- 6 When the **Select which Hard Disk(s) to replicate** page appears, 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 following table for guidance.

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 VMware® vCenter™, appending the replication status to the end.

Option	Description
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. For more information, see Physical backup transport options .
Enable Active Block Mapping (ABM).	<p>Selecting this option enables ABM, which lets vRanger skip data that recognized as removed on Windows® NTFS—including junk data in recycle bin.</p> <p>NOTE: With this option enabled, junk data in recycle bin and deleted data on the VM are not replicated.</p> <p>The following ABM settings options are also available:</p> <ul style="list-style-type: none"> Do not include the Windows page file. Do not include the Windows hibernation file.

- 2 Verify your Transport Failover preference.

If you want to change the Transport Failover order, click **Modify**. For more information, see [Configuring transport failover options](#).

- 3 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.

i | **NOTE:** When using the Full Backup Days option, ensure that the selected recurrence schedule causes backup tasks to 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.

i | **NOTE:** To replicate more frequently than once per day, use the Interval recurrence 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 value 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 value determines the month of the year in which the report occurs.

x can be any value from 1 to 31. This value 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 value determines the month of the year in which the report occurs.
- **Interval:** The Interval option offers the following configurations:
 - **Day(s).** The default is 0, which lets you replicate more frequently than daily.
 - **Hour(s).** The default is 1. To replicate more frequently, change to 0 and configure a **Minute(s)** value.
 - **Minute(s).** The default 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.

i | **TIP:** 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.

i | **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 begin the job automatically after completing the wizard, select **Run the Job after Finish is clicked**.

i | **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 begin the job automatically after clicking Finish, and this job is an on-demand job, do the following:

i | **NOTE:** If this job is not an on-demand job, the job begins at the scheduled time.

- a In the main vRanger UI, 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 On 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**.

Using replication 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 method simplifies the action of failover to a replica VM. You can perform a final synchronization or fail over 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 following describes the default Test Failover operation. To perform a test failover with networking enabled, see [Performing Test Failover with networking enabled](#).

To perform test failure with networking disabled:

- 1 Click the **Test Failover**  icon.
- 2 When the confirmation message appears, 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 When the confirmation message appears, click **Yes**.
The snapshots made are reverted.
The network card is "Connect at Power On," which prepares the replica VM for failover.
The VM returns to its 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](#).

i | **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:

- [Enabling network-enabled Test Failover](#)
- [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. The default is **C:\Program Files\Quest\VRanger**.

To enable network-enabled test failover:

- 1 Open the **Vizioncore.vRanger.Client.Shell.exe.config** file with your text editor of choice.

- 2 Make the required changes carefully, using the following information as guidance.


i | NOTE: Not all text editors include line numbers. If you do not see the following line numbers in your text editor, this absence is not an indication of a problem.

- 3 Scroll to line 42 or search for **Test Failover**.
- 4 Change the value for **ShowTestFailoverDialog** from “false” to “true”, ensuring the quotation marks remain.
- 5 Save and exit the file.
- 6 Restart the Quest 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.
- 2 When the Test Failover dialog box appears, select the preferred option:
 - **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 option 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.

i | CAUTION: This option 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,” which prepares the replica VM for failover.
 - The VM returns to its dormant state.
- 5 When the confirmation prompt appears, click **Yes**.

Performing a failover


If the production site becomes unavailable, or you need to move the workload to the disaster recovery (DR) site, you can perform a failover operation from the vRanger UI.

A failover operation consists of the following actions:

- Turning on the production VM.
- [Optional] Performing a final synchronization between sites. This action 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 (if a hardware failure occurred), data changed since the last replication pass is lost.
- Turn on the DR VM.
- Reverse the replication job, with the DR VM now acting as the source.

i | NOTE: The first replication pass after failing over sends 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.
- 3 When the confirmation message appears, click **Yes**.
- 4 When the **Confirm Synchronization** dialog box appears, select one of the following:
 - **Yes:** This option shuts down the production VM and perform a final synchronization to the DR site. This option 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.

i | NOTE: If you are using differential replication, this option 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 DR replica.

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 DR 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 disaster recovery (DR) VM.
- [Optional] Performing a final synchronization between sites. This action ensures that any data changed on the DR 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 DR VM now acting as the target.
- If the job is enabled, replication continues, sending changes from the source production VM to the DR target.

i | 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.
- 3 When the confirmation message appears, click **Yes**.
- 4 When the **Confirm Synchronization** dialog box appear, select one of the following:

- **Yes:** This option ensures that any data changed on the DR 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 DR 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](#)
- [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 when you create them or you can schedule them to be generated later.

For information on the layout of the My Reports view, see [My Reports view](#).

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 main vRanger UI, click the **My Reports** view.
- 2 Start the Report Wizard by doing one of the following:
 - Right-click one of the nodes, and click **Add**.
 - Select one of the nodes, and then click the **Add** icon on the View toolbar.

- 3 On the **Report Filter** page, enter the information as described in the following table.

Option	Description
Date Range	<p>Narrow report data using these fields:</p> <ul style="list-style-type: none">• Over the last: Enter the number in the first field and select Days or Months from the drop-down list—the default is 30 days.• Between: Click this option 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	<ul style="list-style-type: none">• 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 allows you to filter the report results by repository.
Duration	<p>In this section, you can enter a number in the Greater than or equal to ___ minutes field. The value you enter restricts the report to include only jobs that occurred within that time limit— the default is 0.</p>
Job Status	<p>The check boxes available in this section refer to the status of any job that you might include in reports. Available values are: Completed, Failed, Aborted, and Cancelled.</p> <p>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.</p>

- 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** field 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** field 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** field, 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**.

Body style

Font	Click the browse (...) button to open the Font dialog box; select your preferred Font, Font style, Size, and Effects for the body of the report; and then click OK . NOTE: The default 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.

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 box; 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.

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 box; select your preferred Font, Font style, Size, and Effects for the report footer; and then click OK . NOTE: The footer refers to the 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**.

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 appears, 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.

Option	Description
Recurrence Pattern	<p>In this section, you establish how often the report should be generated. There are five options within this section:</p> <ul style="list-style-type: none"> • Daily: The daily option can be scheduled to run the report every weekday or every <i>x</i> 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: <ul style="list-style-type: none"> ▪ Day <i>[x]</i> of every <i>[y]</i> month: <p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the job occurs.</p> <p><i>y</i> 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.</p> ▪ The <i>[f]</i> <i>[d]</i> of every <i>[y]</i> month(s): <p><i>f</i> can be either: first, second, third, fourth or last.</p> <p><i>d</i> can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p> <p><i>y</i> 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.</p> • Yearly: The yearly option offers the following configurations: <ul style="list-style-type: none"> ▪ Every <i>[m]</i> <i>[x]</i>: <p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p> <p><i>x</i> can be any value from 1 to 31. This value determines the day of the month on which the report occurs.</p> ▪ The <i>[f]</i> <i>[d]</i> of <i>[m]</i>: <p><i>f</i> can be either: first, second, third, fourth, or last.</p> <p><i>d</i> can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.</p> <p><i>m</i> is any month of the year. This value determines the month of the year in which the report occurs.</p> • Interval: The interval option lets you select the number of days, hours, and minutes that should pass between report jobs. <p>NOTE: The interval selected must be greater than or equal to five minutes.</p>
Range of Recurrence	<p>Select the start and end date for running the report.</p> <ul style="list-style-type: none"> • 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**.

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: <ul style="list-style-type: none"> • PDF • Excel

- 9 On the **Report Summary** page, verify that the information is correct, and then click **Finish**.

Integrating and monitoring vRanger

- [Integrating vRanger: overview](#)
- [Monitoring vRanger](#)

Integrating vRanger: overview

Like any component of your virtual infrastructure, vRanger can be used with other applications to achieve more complex tasks.

- [Performing vRanger sweep-to-tape using Quest NetVault Backup](#)
- [Using deduplication with Quest NetVault SmartDisk and vRanger](#)
- [Using Data Domain and vRanger repositories](#)

Performing vRanger sweep-to-tape using Quest NetVault Backup

Sweeping vRanger backups to tape using NetVault Backup lets you combine the vRanger high-speed backup and recovery capabilities with the NetVault Backup 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—for example, `any-jobname-Tape`. The `-Tape` suffix can be changed by editing the Windows 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 jobs 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—for example, 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:

- [Ensuring that system requirements are met](#)
- [Checking the list of prerequisites](#)
- [Preparing the systems](#)

- [Preparing for vRanger sweep-to-tape using NetVault Backup](#)
- [Performing backup operations](#)
- [Performing restore operations](#)

Ensuring that system requirements are met

For vRanger:

- vRanger on Windows Server® 2008 R2 SP1 or later.
- 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, Quest PowerVault LTO standalone and library-based tape devices are ideal for this solution.
- Windows Server 2008 R2 SP1 or later.

Checking the list of prerequisites

This document assumes that both vRanger and NetVault Backup have been installed and configured, and does not cover the installation or the configuration of these applications. For guidance on proper installation and configuration of the items identified in the following list, see the applicable documentation for NetVault Backup and vRanger.

- vRanger 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 tapes 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 Quest PowerVault tape devices, visit our Supported Devices page for a complete list of supported devices.

Preparing the systems

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 from: <https://support.quest.com/netvault-backup/kb/141112>
- Copy the vRanger PowerShell scripts to the NetVault Backup folder.

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 documentation are designed for the first scenario and can be adapted to fit scenario 2.

Scenario 1: vRanger and NetVault Backup are on the same server:

- vRanger resides on the NetVault Backup Server.
- A local repository shared by vRanger and NetVault Backup.
- Windows PowerShell®—PowerShell and vRanger must always be installed on the same server.
- A PowerShell script that finds and starts the vRanger on-demand job.
- A .bat file designed to start the PowerShell script from NetVault Backup Server.

i | NOTE: When installing vRanger and NetVault Backup on the same server, do *not* install the Quest NetVault Backup Plug-in for VMware®.

Scenario 2: vRanger and 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 on the vRanger server, and shared by vRanger and NetVault Backup.
 - Windows PowerShell.
 - A PowerShell script that finds and starts 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:

- [Enabling execution of PowerShell scripts](#)
- [Using vRanger sweep-to-tape scripts](#)
- [Creating a sweep-to-tape vRanger repository](#)

Enabling execution of PowerShell scripts

Windows PowerShell® contains an execution policy that determines whether scripts are allowed to run, and whether they must be digitally signed before they can run. While the default policy is set to block the execution of such scripts, the following topic 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 Click **Start > All Programs > Quest > vRanger Backup & Replication > vRanger Console**.
This step 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 To accept the change, type **Y** and press **Enter**.
- 4 To close the vRanger Console, type **Exit** and press **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, Quest 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** file from: <https://support.quest.com/netvault-backup/kb/141112>
- 2 Extract the contents of the zip file to the following folder on your NetVault Backup Server:

C:\Program Files (x86)\Quest\NetVault Backup\scripts

NetVault Backup only runs 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 Click **Start > All Programs > Quest > vRanger Backup & Replication > vRanger Backup & Replication**.
- 2 In the navigation pane on the left, 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.

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:

- [Creating an on-demand vRanger backup job](#)
- [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 In the navigation pane on the left, select **My Inventory**, 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**.

i | 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® VM, you can customize the method as appropriate on the Transport Selection page.

- 7 On the **Options Selection** page, 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; click **Next**.

i | NOTE: Selecting more savepoints provides additional points-in-time to which you can recover the machine. For more information, see [Understanding retention policies and space-saving technologies](#).

- 9 On the **Recurrence Schedule Selection** page, select **This will be an on demand job**, and 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 click **Next**.
- 11 On the **Summary** page, review the Summary window to make sure that 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. In the following scenario, most of the default settings are used so that the job captures a full backup each time the job runs. For more information on customizing the available NetVault Backup options, see the *Quest NetVault Backup Administrator's Guide*.

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 Start the NetVault Backup WebUI, and click **Create Backup Job**.
- 3 In **Job Name**, specify a name for the job; for example, vRanger_Sweep_to_Tape.

Assign a descriptive name that lets you easily identify the job when monitoring its progress or restoring data. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction; however, a maximum of 40 characters is recommended.
- 4 In the **Selections** list, click **Create New**, 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.
- 5 Navigate to the vRanger repository, and 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.

- 6 In the **Plugin Options** list, select an existing Backup Options Set, or click **Create New**, and configure the options that you want to use.
- 7 In the **Schedule** list, select an existing Schedule Set, or click **Create New**, and configure the schedule type and schedule method.

NetVault Backup provides scheduling options such as immediate, one-time, and triggered.

- 8 In the **Target Storage** list, click **Create New**, and then do the following:
 - a On the **Create Target Set** page, click **Device Selection**.
 - b Select **Specify Device** for a list of available devices.
 - c Select the preferred device in the list to tell NetVault Backup to use this device in the backup.

i | NOTE: Selecting multiple devices lets NetVault Backup use additional devices as required.

- d Select the preferred device in the list to tell NetVault Backup to use this device in the backup.
 - e Click **Set** to save the settings and close the dialog box.
- 9 On the **Create Target Set** page, click **Media Options**, and configure the settings to set specific policies to target specific media, label blank media, or specify whether to reuse media with existing data.
- 10 Next to the **Advanced Options** list, **Create New**, and configure the options that you want to use.

You can select options to customize the backup process; for example, how long to maintain the archive, encryption, and deduplication options; network compression. In this example most of the options remain with their defaults, except for a Pre Script that runs the vRanger On Demand backup.

- 11 On the **Advanced Options** page, click **Pre & Post Scripts**, and select **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.

- 12 To submit the job for scheduling, click **Save & Submit**.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

i | NOTE: A job that uses the Schedule type of **Triggered** is only scheduled when you run the script.

This job runs 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 Navigation pane, click **Create Restore Job**.

On the **Create Restore Job — Choose Saveset** page, the saveset table provides a list of available savesets. The table shows the saveset name (Job Title and Saveset ID), creation date and time, and saveset size.

- 2 Select the saveset that you want to restore, and click **Next**.

i **NOTE:** If the backup index of the selected saveset is offline, the **Confirm load** dialog box is displayed. Click **OK** to load the index from the backup media, and then in the **Load Index** dialog box, type or select the number of days you want to store the index in the NetVault Database.

- 3 On the **Create Selection Set** page, select the archive, for example, vRanger_Sweep_to_Tape, that you want to restore, and click **Next**.

This step restores the entire vRanger repository.

! **CAUTION:** 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 On the **Create Restore Job** page, specify a name for the job.

Assign a descriptive name that lets you easily identify the job when monitoring its progress or restoring data. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction; however, a maximum of 40 characters is recommended.

- 5 In the **Target Client** list, location to restore the file set from the list of available machines.

- 6 In the **Schedule** list, use the default, **Immediate**, to run the job when it is submitted.

- 7 In the **Source Options** list, select the media from which to restore the archive.

- 8 Click **Submit** to run the job.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

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 VM; for example, Repo-001.

vRanger provides options to restore the complete VM, or individual files or folders.

A list of protected VMs display.

- 2 Select the VM containing the data to restore, and then click **File Level Restore** in the top menu to start the process.

A file structure view of the selected VM appears.

- 3 Locate and select the files and folders to restore 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 extracts the files from the vRanger repository for immediate access.

After recovery is complete, you have access to the newly restored files.

Using deduplication with Quest 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.

i | 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 *Quest vRanger Integration Guide for Quest NetVault SmartDisk*.

Using Data Domain and vRanger repositories

vRanger and EMC® Data Domain® appliances provide a simple and efficient method for backing up and recovering VMware® environments. vRanger provides significant enhancements regarding 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 VMs on VMware® ESXi™ hosts.

For more information about integrating vRanger and Data Domain, see the *Quest vRanger Integration Guide for EMC® Data Domain® Boost (DD Boost™)*.

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 the Microsoft 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 to 10999
Replication task failure: logged as an error, event ID 10001
- Virtual Appliance Events: Reserved ID Range 11000 to 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 to 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 Quest vRanger Service, **Vizioncore.vRanger.Service.exe.config**. This file is located in the vRanger installation directory. The default is **C:\Program Files\Quest\vRanger**.

CAUTION: Change only the values listed in the following topics unless directed otherwise by Quest Support.

- [Opening the configuration file](#)
- [Configuring replication monitoring](#)
- [Configuring virtual appliance \(VA\) monitoring](#)
- [Configuring datastore monitoring](#)
- [Restarting the Quest vRanger Service](#)

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** file with your text editor of choice.
- 2 Make the required changes carefully, using the following information as guidance.
 - [Configuring replication monitoring](#)
 - [Configuring virtual appliance \(VA\) monitoring](#)
 - [Configuring datastore monitoring](#)
 - [Restarting the Quest vRanger Service](#)

NOTE: Not all text editors include line numbers. If you do not see the referenced line numbers in your text editor, this absence 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 VM, target host, and vCenter.

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 topic to make additional changes.

Configuring virtual appliance (VA) monitoring

When enabled, at the specified interval, the vRanger VA monitor goes through all the configured VAs and checks to make sure that their power is turned on. The vRanger VA monitor also verifies that vRanger is able to connect to the VAs using the specified IP addresses and credentials.

To configure VA monitoring:

- 1 In the **Vizioncore.vRanger.Service.exe.config** file, find the lines corresponding to the preferred setting:
 - **VirtualApplianceMonitoringEnabled**: line 204.
Set this value to “true” to enable monitoring of the VA.
 - **VirtualApplianceMonitoringEventLogFailureLogMax**: line 207.
This setting 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 next value.
 - **VirtualApplianceMonitorSecondaryIntervallnHours**
This value determines how often the refresher message occurs for a failed VA check after the FailureLogMax value has been reached. The default is **6**, which indicates that it logs an event every six hours.
 - **VirtualAppliancesConnectionCheckIntervallnMinutes**
This setting determines how often vRanger checks the connection to configured VAs. The default is **5**, which indicates that it checks the connection every five minutes.
- 2 Save and exit the file.
Continue to the next topic to make additional changes.

Configuring datastore monitoring

When enabled, at the specified interval, the datastore monitor goes through all the configured datastores and determines if the space available on the datastore exceeds the specified threshold. This check might be either by percentage or a specific value in gigabytes. When this exception 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:
 - **DatastoreMonitoringEnabled**: line 218
Set this value to “true” to enable monitoring of the datastores.
 - **DatastoreMonitoringSpaceThresholdValue**: line 220
This setting is the threshold value for a datastore space before a warning is sent. This value is modified by the **DatastoreMonitoringSpaceThresholdType** setting.
 - **DatastoreMonitoringSpaceThresholdType**: line 222
This value determines whether vRanger checks for percentage of free space or an absolute value in GB. Set “0” for percent, “1” for GB.
 - **DatastoreMonitorCheckIntervallnMinutes**: line 224
How often vRanger checks the datastores for low space (in minutes)—the default is 30 minutes.
 - **DatastoreMonitoringEventLogFailureLogMax**: line 227
The maximum number of times a particular datastore monitor failure can be logged to the Windows event log before a success—the default is **7**.
This setting is intended to limit the number of messages in your event viewer. A refresher message appears periodically based on the next value.
 - **DatastoreMonitorSecondaryIntervallnHours**: line 230

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 is **6**, which indicates it logs an event every six hours.

- 2 Save and exit the file.

Restarting the Quest vRanger Service

For the changes to take effect, you must restart the vRanger Service.

To restart the service:

- 1 On the vRanger machine, click **Start > Run**.
- 2 Type “**services.msc**” and click **OK**.
- 3 Find the vRanger Service. Right-click the service, and click **restart**.

Using the vRanger Console

- [vRanger Console overview](#)
- [Configuring the vRanger Console](#)
- [vRanger vAPI Cmdlets](#)

vRanger Console overview

vRanger uses Windows PowerShell® with a Quest snap-in to provide scripting access to vRanger features and functions. This feature lets vRanger 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 and view relevant information data reading your virtual environment.

For more information on PowerShell, see <https://technet.microsoft.com/library/dn425048.aspx>.

Configuring the vRanger Console

The vRanger installer automatically installs Windows 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 process is a one-time configuration that does not need to be repeated.

To configure the vRanger Console:

- 1 Click **Start > All Programs > Quest > vRanger Backup & Replication > 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 issue is due to the default PowerShell execution policy, which must be enabled.

- 2 At the prompt, type the following command, and press **Enter**.

```
Set-ExecutionPolicy RemoteSigned
```

i | NOTE: PowerShell and vAPI commands are not case-sensitive.

- 3 When the prompt re-appears, close and re-open the vRanger Console.

When the vRanger Console restarts, you should see a screen 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 following cmdlets.

Inventory management

Add-DdbRepository	Get-VirtualApplianceConfig
Add-PhysicalMachine	Get-VirtualApplianceDeploymentStatus
Add-VirtualAppforLinuxFLR	Get-VirtualApplianceReconfigStatus
Add-VirtualAppforLinuxFLRVA	Get-VirtualMachinesUnderInventory
Add-VirtualCenter	Remove-VirtualAppliance
Get-Datastore	Remove-VirtualApplianceConfiguration
Get-InventoryEntity	Update-HypervRestoreJobTemplate
Get-IsInventoryRefreshing	Update-VirtualAppliance
	Update-VirtualApplianceConfiguration

Configuration

Dismount-LinuxVolume	New-TransportConfiguration
Get-MonitorLog	New-VirtualAppliance
Get-VirtualApplianceConfig	Remove-AllMount
Get-VmDisk	Remove-LinuxVolume
Get-VMDKVolume	Set-CBTonVM
Install-VirtualAppliance	Set-LinuxVolume
Mount-LinuxVolume	Set-MountPath
New-EmailAddress	Set-Resources
New-SMTPServer	

Job management

Disable-Job	New-IntervalSchedule
Enable-Job	New-MonthlySchedule
Get-AddressBook	New-WeeklySchedule
Get-ConfigOption	New-YearlySchedule
Get-Connection	Remove-JobTemplate
Get-CurrentTemplateVersionID	Remove-VirtualApplianceConfiguration
Get-Job	Set-Cataloging
Get-JobTemplate	Stop-vRangerJob
New-Daily Schedule	

Repositories

Add-CIFSRepository
Add-DdbRepository
Add-HypervCluster
Add-NVSDRepository
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 preceding vAPI cmdlets.

vAPI Cmdlet details

i **IMPORTANT:** For the most up-to-date help contents, see the cmdlet 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 following 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	<code>Add-BackupGroupEntity [-BackupGroupName] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none"> • 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 backup group: VMware, HyperV, or Physical. If not specified, the default is VMware.
Examples	<code>Add-BackupGroupEntity -BackupGroupName BG1</code>
Related links	Remove-BackupGroupEntity New-BackupGroupMember Remove-BackupGroupMember Get-BackupGroupEntity

Add-BackupJobTemplate

Description	Adds a backup job. You also have the option of running the job immediately after the template is created.
Syntax	<pre>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>]</pre>

Parameters

- **JobName** <String>: Job name.
- **JobDescription** [<String>]: Job description.
- **JobEntity** <Entity>: Inventory entity on which the job is going to run on—VMware® ESXi™ host, VirtualCenter, and so on.
- **ExcludedVMList** [<String>]: List of VM universally unique identifiers (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 store 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 VM 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

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	<ul style="list-style-type: none"> • 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 of 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	<pre>Add-CifsRepository -Server myserver.dev.corp -Sharename myshare -TargetDirectory myrepo -User myname -Password mypassword -Name myCifs Repo -Overwrite 1</pre> <p>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.</p>
Related links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-HypervCluster Add-NVSDRepository

Add-DdbReplicationRepository

Description	Adds an EMC® Data Domain® Boost (DD Boost™) replication repository to vRanger.
Syntax	<pre>Add-DdbReplicationRepository [-Server] <String> [StorageUnit] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>] [-RepositoryId] <String> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • Server <String>: The DD Boost server FQDN or IP. • StorageUnit <String>: The DD Boost storage unit. • User <String>: The user name to access the DD Boost server. • Password <String>: The password to access the DD Boost server. • Name <String>: A friendly name describing this DD Boost server. • Description [<String>]: The description of the DD Boost server. • RepositoryId <String>: Repository ID in GUID. • <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	<pre>Add-DdbReplicationRepository -Server myserver.dev.corp -StorageUnit mystorage -User myname -Password mypassword -Name myDDB Repo -RepositoryId 82E8A98A-1165-4ABE-966B-0E0AEF260620</pre> <p>This command configures a replication repository for repository identified by ID 82E8A98A-1165-4ABE-966B-0E0AEF260620 on server myserver.dev.corp pointing to a storage unit named mystorage using the credentials provided.</p>
Related links	Get-Repository Add-DdbRepository

Add-DdbRepository

Description	Adds an EMC Data Domain Boost (DDB) repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	<pre>Add-DdbRepository [-Server] <String> [-StorageUnit] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-ImportOption] <String>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • Server <String>: The EMC® Data Domain® Boost (DD Boost™) server FQDN or IP. • StorageUnit <String>: The DD Boost storage unit. • User <String>: The user name to access the DD Boost server. • Password <String>: The password to access the DD Boost server. • Name <String>: A friendly name describing this DD Boost server. • Description [<String>]: The description of the DD Boost 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	<pre>Add-DdbRepository -Server myserver.dev.corp -StorageUnit mystorage -User myname -Password mypassword -Name myDDB Repo -Overwrite 1</pre> <p>This command adds a repository on server myserver.dev.corp pointing to a storage unit named mystorage using the credentials provided. Existing repositories at the location specified are overwritten.</p>
Related links	Get-PhysicalMachineDiskMap Add-CIFSRepository Add-HypervCluster Add-NVSDRepository

Add-EsxHost

Description	Adds a VMware® ESXi™ host to the inventory.
Syntax	<code>Add-EsxHost [-DNSorIP] <String> [-Username] <String> [-RootPassword] <String> [[-UserPassword] <String>] [[-Port] <Int>] [[-AddUser] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• DNSorIP <String>: The DNS or IP for the VMware® ESXi™ host.• Username <String>: The user name for the ESXi host.• RootPassword <String>: The root password for the ESXi host.• UserPassword [<String>]: The user password for the ESXi host.• Port [<Int>]: The SSH Port for the ESXi host.• AddUser [<Boolean>]: Add user if 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	<pre>Add-EsxHost -DNSorIP 10.6.20.20 -Username root -RootPassword p@ssw0rd</pre> <p>This command adds an ESXi host to the inventory.</p>
Related links	Add-VirtualCenter Get-InventoryEntity

Add-HypervCluster

Description	Adds a Hyper-V® Cluster to the inventory.
Syntax	<code>Add-HypervCluster [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int32>] [[-IsInstallAgent] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• DNSorIP <String>: The DNS or IP for the Hyper-V cluster.• Username <String>: The user name 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	<pre>Add-HypervCluster -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd</pre> <p>This command adds a Hyper-V cluster to the inventory.</p>
Related links	Add-BackupJobTemplate Get-InventoryEntity

Add-HypervHost

Description	Adds a Hyper-V® host to the inventory.
Syntax	<code>Add-HypervHost [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int32>] [[-IsInstallAgent] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• DNSorIP <String>: The DNS or IP for the Hyper-V host.• Username <String>: The user name 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	<pre>Add-HypervHost -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd</pre> <p>This command adds a Hyper-V host to the inventory.</p>
Related links	Add-BackupJobTemplate Get-InventoryEntity

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	<code>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>]</code>

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*>]: Separate multiple email addresses with commas.
- **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 for the host is used.
- **SkipDisks** [<*List`1*>]: Sequence number—base 0—of the disks to skip from the savepoint. Separate multiple numbers with commas.
- **DiskLocations** [<*List`1*>]: Disk locations to use on the target host. Separate multiple values with commas. 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. Separate multiple values with 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. Separate multiple numbers with 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 '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 restores 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	Adds a Network File Share repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	<pre>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>]</pre>
Parameters	<ul style="list-style-type: none">• 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	<p>-----EXAMPLE 1-----</p> <pre>Add-NfsRepository -Server myserver.dev.corp -ExportDirectory /vol/vol1 -TargetDirectory /test -Name MyNFSRepo -Overwrite 1</pre> <p>This command adds a repository on server myserver.dev.corp. Existing repositories at the location specified are overwritten.</p> <p>-----EXAMPLE 2-----</p> <pre>Add-NfsRepository -Server myserver.dev.corp -ExportDirectory /vol/vol1 -TargetDirectory /test -Name MyNFSRepo -EnableEncryption 1 -EncryptionPassword myPassword -EnableCompress 1</pre> <p>This command adds a repository on server myserver.dev.corp. Enables encryption on backups of this repository. All backups to this repository are compressed.</p>
Related links	Get-PhysicalMachineDiskMap Add-DdbRepository Add-NVSDRepository Add-CIFSRepository

Add-NVSDRepository

Description	Adds a Quest NetVault SmartDisk repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	<pre>Add-NvsdRepository [-Server] <String> [-TargetDirectory] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-ImportOption] <String>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none">• 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	<pre>Add-NvsdRepository -Server myserver.dev.corp -TargetDirectory myrepo -User myname -Password mypassword -Name myNVSD Repo -Overwrite 1</pre> <p>This command adds a repository on server myserver.dev.corp pointing to target directory named myrepo using the credentials provided. Existing repositories at the location specified will be overwritten.</p>
Related links	Get-Repository Add-NFSRepository Add-DdbRepository Add-CIFSRepository Add-RdaRepository

Add-PhysicalMachine

Description	Adds a physical machine to the inventory.
Syntax	<pre>Add-PhysicalMachine [-DNSorIP] <String> [-Username] <String> [-Password] <String> [[-Port] <Int>] [-IsInstallAgent] <Boolean> [[-AgentLocation] <String>] [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • DNSorIP <String>: The DNS or IP for the physical machine. • Username <String>: The user name 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	<pre>Add-PhysicalMachine -DNSorIP 10.6.20.20 -Username user1 -Password p@ssw0rd</pre> <p>This command adds a replication job template.</p>
Related links	Add-BackupGroupEntity Get-InventoryEntity

Add-RdaRepository

Description	Adds an RDA repository to vRanger. The repository can be later used to store backup files from backup jobs.
Syntax	<pre>Add-RdaRepository [-Server] <String> [-StorageUnit] <String> [-User] <String> [-Password] <String> [-Name] <String> [[-Description] <String>] [[-Overwrite] <Boolean>] [[-ImportOption] <String>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • Server <String>: The DR Series server FQDN or IP. • StorageUnit [<String>]: The RDA logical storage unit. • User <String>: The user name to access the DR Series server. • Password <String>: The Password to access the DR Series server. • Name <String>: A friendly name describing this DR Series 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	<pre>Add-RdaRepository -Server myserver.dev.corp -StorageUnit mystorage -User myname -Password mypassword -Name myRDAREpo -Overwrite 1</pre> <p>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 will be overwritten.</p>
Related links	Get-Repository Add-DdbRepository Add-CIFSRepository Add-NVSDRepository

Add-ReplicationJobTemplate

Description	Adds a replication job template. You also have the option of running the job immediately after the template is created.
Syntax	<pre>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>]</pre>
Parameters	<ul style="list-style-type: none"> • 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>]: Name of replicated VM. • Type <String>: Replication type. Valid types are Differential and Hybrid. • TargetNetworks [<String>]: Networks to use on the target host. Separate multiple values with commas. • PrimaryDatastore <String>: Primary Datastore on target host where the replicated VM is to be located. • TargetDatastores [<String>]: Datastores to use on the target host. Separate multiple values with commas. • JobSchedule <Schedule>: The recurrence schedule on which the job runs. • TransportFailover [<String>]: Transport Failover order separated by commas to be attempted until the task is successful. Valid types are VA-HotAdd and VA-Network. • NotificationList [<String>]: Email notification list for the job. Separate multiple email addresses with commas. • Flags <ReplicationFlags>: Replication flags: NoFlags, ReplicatePoweredOnOnly, CheckDestinationFreeSpace, UseCompression, UpdateNotes, EnableGuestQuiescing, AllFlags. Separate multiple flags with pipes (). • 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 VM Network -
includeddisk 0,1,2
```

This command adds a replication job template. Since the RunJobNow flag is not specified the job will not be run immediately. Since no transport failover is specified the default order of transport failovers will be used.

-----EXAMPLE 2-----

```
Add-ReplicationJobTemplate -JobName vRepJob1 -JobEntity $vm -
TargetHost $targethost -ReplicateName TestRep -Type Differential -
Flags $flags -targetDatastores MyDatastore1, MyDatastore2 -
PrimaryDatastore MyDatastore -targetnetworks VM Network -
TransportFailover VA-HotAdd, VA-Network -RunJobNow 1 -includeddisk
0,2,4
```

This command adds a replication job template.

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. You 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.
- **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>]**: Separate multiple email addresses with commas.
- **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. Separate multiple numbers with commas.
- **TargetDatastores** **[<String>]**: Datastores to use on the target host. Separate multiple values with commas. 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. Separate multiple values with 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; 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. Separate multiple numbers with commas.
- **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"

-----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 an 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. You also have the option of running the job immediately after the template is created.
Syntax	<pre>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>]</pre>

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. Separate multiple flags with pipes (|).
- **NotificationList** **[<String>]**: Email notification list for the job. Separate multiple email addresses with commas.
- **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. Separate multiple numbers with commas.
- **TargetDatastores** **[<String>]**: Datastores to use on the target host. Separate multiple values with commas. 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. Separate multiple values with 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; 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. Separate multiple numbers with commas.
- **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"

-----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 restores 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 restores 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 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'}
$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-VirtualAppforLinuxFLR

Description	This command is obsolete and has been replaced with Add-VirtualAppforLinuxFLRVA .
Related links	Add-VirtualAppforLinuxFLRVA

Add-VirtualAppforLinuxFLRVA

Description	Adds a VA to perform Linux® FLR.
Syntax	<code>Add-VirtualAppforLinuxFLRVA [-VirtualApplianceUUID] <String> [<i><CommonParameters></i>]</code>
Parameters	<ul style="list-style-type: none">• VirtualApplianceUUID <String>: VA UUID.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>Add-VirtualAppforLinuxFLRVA -VirtualApplianceUUID 564d8ab2-6c7f-5915-0a94-b224b5a71570</code>
Related links	None

Add-VirtualCenter

Description	Adds a VirtualCenter to the inventory.
Syntax	<code>Add-VirtualCenter [-DNSorIP] <String> [-Username] <String> [-UserPassword] <String> [[-Port] [<Int>]] [<i><CommonParameters></i>]</code>
Parameters	<ul style="list-style-type: none">• DNSorIP <String>: The DNS or IP for the VirtualCenter.• Username <String>: The user name for the VirtualCenter.• UserPassword <String>: The 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	<code>Add-VirtualCenter -DNSorIP 10.6.20.20 -Username administrator -RootPassword p@ssw0rd</code> This command adds a VirtualCenter to the inventory.
Related links	Add-DdbRepository Get-InventoryEntity

Disable-Job

Description	Disables enabled jobs.
Syntax	<code>Disable-Job [-JobTemplateIDs] <GUID[]> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<p>-----EXAMPLE 1-----</p> <pre>Disable-Job -JobTemplateId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'</pre> <p>This command disables the job if it is enabled and returns the job template entity.</p> <p>-----EXAMPLE 2-----</p> <pre>\$jobs = Get-JobTemplate where {\$_.IsEnabled -eq 1} Disable-Job \$jobs[0].Id</pre> <p>This command disables the first enabled job in the list and returns the job template entity.</p>
Related links	Enable-Job Remove-VirtualApplianceConfiguration Get-JobTemplate Remove-JobTemplate

Dismount-LinuxVolume

Description	Unmounts VAR disk from the Linux® FLR VA.
Syntax	<code>Dismount-LinuxVolume [<CommonParameters>]</code>
Parameters	<CommonParameters> : This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>Dismount-LinuxVolume</code>
Related links	Mount-LinuxVolume

Enable-Job

Description	Enables disabled jobs.
Syntax	<code>Enable-Job [-JobTemplateIDs] <GUID[]> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<p>-----EXAMPLE 1-----</p> <pre>Enable-Job -JobTemplateId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'</pre> <p>This command enables the job if it is disabled and returns the job template entity.</p> <p>-----EXAMPLE 2-----</p> <pre>\$jobs = Get-JobTemplate where {\$_.IsEnabled -eq 0} Enable-Job \$jobs[0].Id</pre> <p>This command enables the first disabled job in the list and returns the job template entity.</p>
	<p>Related links</p> <ul style="list-style-type: none"> Disable-Job Remove-VirtualApplianceConfiguration Get-JobTemplate Remove-JobTemplate

Get-AddressBook

Description	Gets a list of email addresses from the address book.
Syntax	<code>Get-Addressbook [<CommonParameters>]</code>
Parameters	<CommonParameters> : This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>Get-Addressbook</code>
Related links	New-EmailAddress

Get-BackupGroupEntity

Description	Returns a list of backup group entities.
Syntax	<code>Get-BackupGroupEntity [<CommonParameters>]</code>
Parameters	<CommonParameters> : This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<p>-----EXAMPLE 1-----</p> <pre>Get-BackupGroupEntity</pre> <p>This command gets a list of all backup group entities.</p> <p>-----EXAMPLE 2-----</p> <pre>Get-BackupGroupEntity where{\$_.Name -eq 'BackupGroupName'}</pre> <p>This command gets a list of backup groups by name.</p>
	None
Related links	None

Get-CatalogSearchData

Description	Returns the list of results for the matched search string, repositories, and VMs.
Syntax	<code>Get-CatalogSearchData [-SearchText] <String> [[-RepositoryList] [<String>]] [[-VMList] [-Type] <Entity>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<p>-----EXAMPLE 1-----</p> <pre>Get-CatalogSearchdata -SearchText abcd</pre> <p>Gets the result set from the catalog database for matched search text. No repository or machine filtering is applied.</p> <p>-----EXAMPLE 2-----</p> <pre>\$repo = Get-Repository where {\$_.Name -eq 'Repo1'} \$vmList = Get-InventoryEntities -Type VirtualMachine {\$_.Name -eq 'VM1'} Get-CatalogSearchdata -SearchText abcd -RepositoryList \$repo -VMList \$vm</pre> <p>Gets the result set from the catalog database for matched search text file read by repository and machine.</p>
Related links	Remove-Catalog Get-CatalogStatus Set-Cataloging

Get-CatalogStatus

Description	Gets the active status of Catalog for a list of savepoints.
Syntax	<code>Get-CatalogStatus [-Savepoints] <GUID[]> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>Get-CatalogStatus -Savepoints 5a62765e-6a04-4df4-af54-46e5a38ad876</pre> <p>Gets the active status of the catalog data for the associated savepoint.</p>
Related links	Get-CatalogSearchData Remove-Catalog Set-Cataloging

Get-ConfigOption

Description	Returns the config option values for a specified type.
Syntax	<code>Get-Configoption [-Type] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• Type <String>: Valid types are Database, Email, Resources, Timeout, General, LinuxFLR, and Cataloging.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>Get-Configoption -Type Resources</code>
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	<code>Get-Connection [[-Name] [<String>]] [[-Type] [<String>]] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• Name [<String>]: The name or IP address of the VMware® ESXi™ host, VMware® 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	<p>-----EXAMPLE 1-----</p> <pre>Get-Connection -type HypervHost</pre> <p>This command gets a list of all Hyper-V host connections.</p> <p>-----EXAMPLE 2-----</p> <pre>Get-Connection</pre> <p>This command gets a list of all vRanger loaded connections.</p> <p>-----EXAMPLE 3-----</p> <pre>Get-Connection -Name myhost.dev.corp</pre> <p>This command returns a connection to the specified host.</p> <p>-----EXAMPLE 4-----</p> <pre>Get-Connection -type EsxHost</pre> <p>This command gets a list of all ESXi host connections.</p>
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	<code>Get-CurrentTemplateVersionID [-TemplateVersionID] [<String>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<code>Get-CurrentTemplateVersionID 6baeede2-227a-4350-ba3f-470f9e90dc07</code> 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	Returns all datastores associated with the specified host.
Syntax	<code>Get-Datastore [-Host] <HostEntity> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<code>Get-Datastore -Host \$host</code> 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	<code>Get-GlobalTransportFailover [<CommonParameters>]</code>
Parameters	<CommonParameters> : This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>Get-GlobalTransportFailover</code> This command gets the default Transport Failover order for replication jobs.
Related links	Update-GlobalTransportFailover Get-TransportFailover

Get-InventoryEntity

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. There is a performance delay if all types are requested or no type is specified.
Syntax	<code>Get-InventoryEntity [[-Type] <String>] [[-Block] <Boolean>] [[-Thin] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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 is done. Possible timeout might 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	<p>-----EXAMPLE 1-----</p> <pre>Get-InventoryEntity- Type HypervHost</pre> <p>This command gets a list of all inventory entities.</p> <p>-----EXAMPLE 2-----</p> <pre>Get-InventoryEntity -Type EsxHost</pre> <p>This command gets a list of all loaded VMware® ESXi™ host entities.</p> <p>-----EXAMPLE 3-----</p> <pre>Get-InventoryEntity</pre> <p>This command gets a list of all inventory entities.</p>
Related links	None

Get-IsInventoryRefreshing

Description	Returns "true" or "false" based on whether the inventory is refreshing at the time the cmdlet is run.
Syntax	<code>Get-IsInventoryRefreshing [[-Type] <String>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	None

Get-Job

Description	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	<ul style="list-style-type: none"> • 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 with commas. • <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 -JobTemplateIDs d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3 This command returns jobs based on the specified template ID.
Related links	Get-JobTemplate

Get-JobTemplate

Description	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	<pre>Get-JobTemplate [[-Type] <String>] [[-InventoryIds] <String[]>] [[-Thin] <Boolean>] [[-ID] <String>] [-VersionID] <String>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none">• 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 with commas.• Thin [<Boolean>]: Whether to return a 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	<p>-----EXAMPLE 1-----</p> <pre>Get-JobTemplate</pre> <p>This command gets a list of all job templates.</p> <p>-----EXAMPLE 2-----</p> <pre>Get-JobTemplate -type Backup</pre> <p>This command gets a list of all backup jobs.</p> <p>-----EXAMPLE 3-----</p> <pre>Get-JobTemplate -InventoryIds d7ebc5c4-7916-44c4-a0c2- 2412999735f3,d7ebc5c4-7916-44c4-a0c2-24 12990005f3 -Type Backup</pre> <p>This command returns backup job templates in the specified inventories.</p>
Related links	Get-Job

Get-MonitorLog

Description	Displays the vRanger Windows® error events.
Syntax	<pre>Get-MonitorLog [[-EventId] <Int32>] [[-StartDateTime] <DateTime>] [[-EndDateTime] <DateTime>] [-Verbose] [-Debug] [-ErrorAction <ActionPreference>] [-WarningAction <ActionPreference>] [-ErrorVariable <String>] [-WarningVariable <String>] [-OutVariable <String>] [-OutBuffer <Int32>] [-WhatIf] [-Confirm]</pre>
Parameters	<ul style="list-style-type: none">• EventId <Int32>: Filter by event ID.• StartDateTime <DateTime>: Filter based on the start time.• EndDateTime <DateTime>: Filter based on the end time.

Examples	<code>get-monitorlog -eventid 1000</code>
Related links	None

Get-Network

Description	Returns all networks associated with the host specified.
Syntax	<code>Get-Network [-Host] <HostEntity> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none"> • 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	<p><code>Get-Network -Host \$host</code></p> <p>This command returns a list of networks associated with the host.</p>
Related links	Get-Datastore

Get-PhysicalMachineDiskMap

Description	Creates a Physical Machine disk map to be used in AddRestoreJobTemplate.
Syntax	<code>Get-PhysicalMachineDiskMap [-RestoreCDIP] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none"> • RestoreCDIP <String>: The physical 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	<p><code>GetPhysicalMachineDiskMap -RestoreCDIP \$IP</code></p> <p>Creates a Physical Machine disk map to be used in AddRestoreJobTemplate.</p>
Related links	Get-InventoryEntity 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 parameters in the same command returns an error.
Syntax	<code>Get-Repository [[-Id] [<String>]] [[-Type] <String>] [<CommonParameters>]</code>

Parameters	<ul style="list-style-type: none"> • Type <String>: Specifies the type of repository to retrieve: CIFS, 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	<p>-----EXAMPLE 1-----</p> <pre>Get-Repository</pre> <p>This command gets a list of all vRanger repositories.</p> <p>-----EXAMPLE 2-----</p> <pre>Get-Repository -id 034538e5-58a7-4d28-89c9-3249b620064a</pre> <p>This command a repository with the specified ID.</p> <p>-----EXAMPLE 3-----</p> <pre>Get-Repository -type CIFS</pre> <p>This command gets a list of all CIFS repositories.</p>
Related links	Add-CIFSRepository Add-DdbRepository Add-HypervCluster Remove-Repository

Get-RepositoryJob

Description	Returns a list of associated jobs based on the input repository ID.
Syntax	<code>Get-RepositoryJob [-RepositoryId] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-RepositoryJob -RepositoryId 034538e5-58a7-4d28-89c9-3249b620064a</pre> <p>This command returns a list of jobs associated with the repository with the specified ID.</p>
Related links	Get-Job Get-JobTemplate

Get-RepositorySavePoint

Description	Returns a list of savepoints in the repository. All savepoints in the repository are returned if no time parameter is specified.
Syntax	<code>Get-RepositorySavePoint [-Id] <String> [[-StartTime] [<String>]] [[-EndTime] [<String>]] [<CommonParameters>]</code>

Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-RepositorySavePoint -Id 034538e5-58a7-4d28-89c9-3249b620064a</pre> <p>This command returns a list of savepoints associated with the repository with the specified ID.</p>
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	<pre>Get-RestoreDiskMap [-PhysicalDiskMap] <PhysicalDiskMap> [-ManifestFile] <ManifestFile> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-RestoreDiskMap -PhysicalDiskMap \$physicaldiskMap -ManifestFile \$manifestFile</pre> <p>This command matches the file map entries to the disks in a physical machine by size and returns a hash map that you can manipulate by code.</p>
Related links	Add-BackupJobTemplate Get-InventoryEntity

Get-SavepointDisk

Description	Returns the VM disks that the savepoint contains.
Syntax	<pre>Get-SavepointDisk [-Id] <String> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-SavePointDisk 91d27db5-4562-4456-918d-b0dc587e0e7f</pre> <p>This command returns all disks that have been backed up into the savepoint specified.</p>
Related links	Get-VMDKVolume

Get-SavepointManifest

Description	Returns manifest of the savepoint.
Syntax	<pre>Get-SavePointManifest [-ID] <ID> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> 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	<pre>Get-SavePointManifest \$id</pre> <p>This command returns the savepoint manifest object.</p>
Related links	Get-RepositorySavePoint Get-VMDKVolume

Get-Savepoints

Description	Returns a list of savepoints for a specified BackupJobTemplateID.
Syntax	<pre>Get-SavePoints [-BackupJobTemplateID] <GUID> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> 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	<pre>Get-SavePoints -BackupJobTemplateID 034538e2-58a7-4d28-89c9-3249b620064a</pre> <p>This command returns a list of savepoints associated with the backup template with the specified ID.</p>
Related links	Get-Job Get-JobTemplate

Get-TransportFailover

Description	Gets the Transport Failover order for replication jobs by TemplateVersion ID.
Syntax	<pre>Get-TransportFailover [-TemplateVersionID] <String> [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • 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	<pre>\$jobtemplates = Get-JobTemplate -Type Replication \$template = \$jobtemplates[0] Get-TransportFailover -TemplateVersionID \$template.TemplateVersionId</pre> <p>Gets the Transport Failover order for replication jobs by TemplateVersion ID.</p>
Related links	Get-GlobalTransportFailover Update-GlobalTransportFailover

Get-VirtualApplianceConfig

Description	Gets the VirtualApplianceConfig configuration from vRanger.
Syntax	<pre>Get-VirtualApplianceConfig [[-VName] [<String>]] [[-HostName] [<String>]] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • VName [<String>]: The VA Name. • HostName [<String>]: The host name where the VA resides. • <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<p>-----EXAMPLE 1-----</p> <pre>\$va = Get-VirtualApplianceConfig</pre> <p>This command gets all the VA entities.</p> <p>-----EXAMPLE 2-----</p> <pre>\$va = Get-VirtualApplianceConfig -VName ABC-Vzva-1.1.1</pre> <p>This command gets the VA entity for the VA specified.</p> <p>-----EXAMPLE 3-----</p> <pre>\$va = Get-VirtualApplianceConfig -HostName xyz.dev.eng.int</pre> <p>This command gets the VA entities for the Host specified.</p>
Related links	New-VirtualAppliance Update-VirtualAppliance Remove-VirtualAppliance

Get-VirtualApplianceDeploymentStatus

Description	Gets the status of the re-configuration of VA in progress.
Syntax	<pre>Get-VirtualApplianceDeploymentStatus [-DeploymentId] <GUID> [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-VirtualApplianceReconfigStatus -Id \$deploymentID</pre> <p>This command gets the status of the VA re-configuration.</p>
Related links	Update-VirtualAppliance New-VirtualAppliance

Get-VirtualApplianceReconfigStatus

Description	Gets the status of the deployment of VA in progress.
Syntax	<pre>Get-VirtualApplianceDeploymentStatus [-DeploymentId] <GUID> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Get-VirtualApplianceDeploymentStatus -Id \$deploymentID</pre> <p>This command gets the status of the VA deployment in progress.</p>
Related links	Update-VirtualAppliance New-VirtualAppliance

Get-VirtualMachinesUnderInventory

Description	Gets the child VM entities of the entity—host, cluster, and so on—provided.
Syntax	<pre>Get-VirtualMachinesUnderInventory [[-InventoryEntityMoRefId] [<String>]] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • InventoryEntityMoRefId [<String>]: The MoRef ID of the parent inventory entity. • 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	<pre>Get-VirtualMachinesUnderInventoryEntity -InventoryEntityMoRefId host-16641</pre> <p>The managed object reference ID of an inventory entity node such as a VMware® ESXi™ host or resource folder.</p>
Related links	Add-BackupGroupEntity Get-BackupGroupEntity Get-InventoryEntity

Get-VmDisk

Description	Returns a list of VmDisks for the VM list in the same order of the input.
Syntax	<code>Get-VmDisk [-Vmlist] <VMEntity[]> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>Get-VmDisk -Vmlist \$mylist</pre> <p>This command gets a list of VMDisks objects for the input VM list.</p>
Related links	None

Get-VMDKVolume

Description	Returns all volumes on a certain VM disk.
Syntax	<code>Get-VMDKVolume [-VmdkID] <String> [-Repository] <Repository> [-RepositoryPassword] <String> [SavepointManifest] <ManifestFile> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>Get-VMDKVolume -vmdkid \$diskid -repository \$repo -repositorypassword Password -savepointmanifest \$manifest</pre> <p>This command returns all volumes of the specified VMDK.</p>
Related links	None

Install-VirtualAppliance

Description	Deploys a new VA for vRanger.
Syntax	<pre>Install-VirtualAppliance [-VName] <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>]</pre>

Parameters

- **VAName** **<String>**: The VA name.
- **SelectedHost** **<HostEntity>**: The host to associate the VA with.
- **OVFLocation** **<String>**: The path to the local OVF.
- **DownloadOVF** **<Boolean>**: Should the OVF be downloaded?
- **IPAddressOverride** **[<String>]**: IP Address to use for the VA instead of the one retrieved from the VirtualCenter API.
- **ClusterDefault** **[<Boolean>]**: If the selected host belongs to a cluster, share the VA among all hosts in the cluster.
- **CPUAllocation** **<Int>**: Allocate CPU.
- **MemorySizeInMB** **<Int>**: Assign memory in MB.
- **VADatastore** **<Datastore>**: Select a datastore for the VA.
- **NetworkEntity** **<Network>**: Select a network for the VA.
- **ScratchDiskDatastore** **[<Datastore>]**: Select a datastore for the scratch disk.
- **ScratchDiskSizeInKB** **[<Int>]**: Assign scratch disk in KB.
- **NewPassword** **[<String>]**: New password for the VA; if this word is not mentioned, Original password (vzroot1) is used.
- **UseVAForFLR** **[<Boolean>]**: Use VA for Linux® FLR.
- **PowerOnVA** **[<Boolean>]**: Power on VA 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>]**: Netmask 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. For more information, see the Linux manual (man) page for resolv.conf.
- **Search** **[<String>]**: Optional Search list for host name lookup for DNS on VA. For more information, see the Linux man page for resolv.conf.
- **<CommonParameters>**: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"

Examples	<pre> \$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 -VName 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 </pre> <p>This command deploys a new VA configuration to vRanger using a locally saved OVF and retrieves the status of the currently deploying VA in progress. Default for VADeployStatus is false.</p>
Related links	Get-VirtualApplianceConfig Remove-VirtualAppliance Update-VirtualAppliance

Mount-LinuxVolume

Description	Mounts a volume from the specified manifest FileMaps.
Syntax	<pre> Mount-LinuxVolume [-ManifestFileMaps] <FileMapEntryLine[]> [-Repository] <Repository> [-RepositoryPassword] <String> [<CommonParameters>] </pre>
Parameters	<ul style="list-style-type: none"> 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	<pre> \$path = Mount-LinuxVolume -repositorypassword xxxxxxxx -manifestfilemap [filemap] -repository [repo] </pre>
Related links	Dismount-LinuxVolume

New-BackupFlag

Description	Creates a new backup flag for later use.
Syntax	<pre> 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>] </pre>

Parameters	<ul style="list-style-type: none"> • BackupPoweredOnOnly [<Boolean>]: Whether to back up only VMs 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 ABM 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	<p>-----EXAMPLE 1-----</p> <pre>New-BackupFlag</pre> <p>This command returns a new BackupFlags object with the value of NoFlags.</p> <p>-----EXAMPLE 2-----</p> <pre>New-BackupFlag -EnableGuestQuiescing 1 -EnableABM 1</pre> <p>This command returns a new BackupFlags object with the values of EnableGuestQuiescing and EnableABM.</p> <p>-----EXAMPLE 3-----</p> <pre>New-BackupFlag -AllFlags 1</pre> <p>This command returns a new BackupFlags object with all flags set to on.</p>
Related links	Add-BackupJobTemplate

New-BackupGroupMember

Description	Adds members to a backup group created using Add-BackupGroupEntity cmdlet or UI. You can add any inventory item as a member.
Syntax	<pre>New-BackupGroupMember [-BackupGroupId] <GUID> [-MemberEntity] <Entity> [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • BackupGroupId <GUID>: The backup group 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	<pre>\$inventoryEntity = Get-InventoryEntity -type virtualmachine where {\$_ .Name -eq 'VMNAMEHERE'}</pre> <pre>New-BackupGroupMember 8c40f593-6608-45ef-b2c9-a2594aac2338 \$inventoryEntity</pre>
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	<pre>New-DailySchedule [-WeekdayOnly] <Boolean> [-StartTime] <String> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>New-DailySchedule -WeekdayOnly 1 -StartTime "2009-03-31 12:00:00AM"</pre> <p>This command returns a daily job schedule that runs weekdays at midnight.</p>
Related links	None

New-EmailAddress

Description	Adds an email address to the address book.
Syntax	<pre>New-EmailAddress [-Name] <String> [-Email] <String> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>New-EmailAddress -Name Employee1 -Email firstname.lastname@company.com</pre>
Related links	Get-AddressBook

New-IntervalSchedule

Description	Returns an interval schedule object.
Syntax	<pre>New-IntervalSchedule [[-IntervalDays] <Int>] [[-IntervalHours] <Int>] [-IntervalMinutes] <Int> [-StartTime] <String> [-RangeEndDate] <String> [-RangeByOccurrence] <Int>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none">• IntervalDays <Int>: Number of days to use for interval schedule—the default is 0.• IntervalHours <Int>: Number of hours to use for interval schedule—the default is 0.• IntervalMinutes <Int>: Number of minutes to use for interval schedule—the default is 0.• StartTime <String>: Job start time.• RangeEndDate <String>: Range end date.• RangeByOccurrence <Int>: The maximum number of times a job is run.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<p>-----EXAMPLE 1-----</p> <pre>New-IntervalSchedule -StartTime 2009-03-31 -IntervalHours 1</pre> <p>This command returns an interval job schedule that runs every hour.</p> <p>-----EXAMPLE 2-----</p> <pre>New-IntervalSchedule -StartTime 2009-03-31 -IntervalMinutes 30 -RangeEndDate 2010-12-31</pre> <p>This command returns an interval job schedule that runs every 30 minutes until 2010-12-31.</p> <p>-----EXAMPLE 3-----</p> <pre>New-IntervalSchedule -StartTime 2009-03-31 -IntervalHours 1 -intervalMinutes 30 -RangeByOccurrence10</pre> <p>This command returns an interval job schedule that runs every hour and half; the job is run 10 times at this interval.</p>
Related links	New-Daily Schedule New-WeeklySchedule New-MonthlySchedule New-YearlySchedule

New-MonthlySchedule

Description	Returns a Schedule object. The schedule provides a template monthly schedule that can be further customized before use.
Syntax	<pre>New-MonthlySchedule [-ExecutionDayOfMonth] <Int> [-StartTime] <String> [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • ExecutionDayOfMonth <i><Int></i>: Day of month on which the job runs—1 to 31. • StartTime <i><String></i>: 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	<pre>New-MonthlySchedule -ExecutionDayOfMonth 1 -StartTime "2009-03-31 12:00:00AM"</pre> <p>This command returns weekly job schedule that runs on the first of every month at midnight.</p>
Related links	New-Daily Schedule New-WeeklySchedule New-YearlySchedule New-IntervalSchedule

New-ReplicationFlag

Description	Helps you 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	<pre>New-ReplicationFlag [[-CheckDestinationFreeSpace] [<Boolean>]] [[-UpdateNotes] [<Boolean>]] [[-EnableGuestQuiescing] [<Boolean>]] [[-EnableABM] [<Boolean>]] [[-AllFlags] [<Boolean>]] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • CheckDestinationFreeSpace <i><Boolean></i>: Whether to check repository free space before replication. • UpdateNotes <i><Boolean></i>: Whether to update notes on VM after replication. • EnableGuestQuiescing <i><Boolean></i>: Whether guest quiescing is enabled. • EnableABM <i><Boolean></i>: Whether ABM is enabled. • AllFlags <i><Boolean></i>: 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	<p>-----EXAMPLE 1-----</p> <pre>NewReplicationFlag</pre> <p>This command returns a new ReplicationFlags object with the value of NoFlags.</p> <p>-----EXAMPLE 2-----</p> <pre>New-ReplicationFlag -EnableGuestQuiescing 1 -EnableABM 1</pre> <p>This command returns a new ReplicationFlags object with the values of EnableGuestQuiescing and EnableABM.</p> <p>-----EXAMPLE 3-----</p> <pre>New-ReplicationFlag -AllFlags 1</pre> <p>This command returns a new ReplicationFlags object with all flags set to on.</p>
Related links	Add-PhysicalMachine

New-RestoreFlag

Description	Helps you 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	<pre>New-RestoreFlag [[-ForceOverwrite] <Boolean>] [[-ForcePowerOn] <Boolean>] [[-UseFiber] <Boolean>] [[-PerformNetworkRestoreOnFailure] <Boolean>] [[-LanFree] <Boolean>] [[-AllFlags] <Boolean>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none">• 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® VM 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	<p>-----EXAMPLE 1-----</p> <pre>New-RestoreFlag</pre> <p>This command runs RestoreFlags with the value of NoFlags.</p> <p>-----EXAMPLE 2-----</p> <pre>New-RestoreFlag -ForcePowerOn 1 -UseCompression 0</pre> <p>This command runs RestoreFlags with the values of ForcePowerOn.</p> <p>-----EXAMPLE 3-----</p> <pre>New-RestoreFlag -AllFlags 1</pre> <p>This command runs RestoreFlags with all flags set to on.</p>
Related links	None

New-SMTPServer

Description	Adds an SMTP Server.
Syntax	<pre>New-SMTPServer [-SMTPServer] <String> [-Port] <Int> [-FromEmail] <String> [<CommonParameters>]</pre>

Parameters	<ul style="list-style-type: none"> • SMTPServer <String>: SMTP server address. • Port <Int>: Port number for SMTP server. • FromEmail <String>: Email address for sent email messages. • <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<code>New-SMTPServer -SMTPServer company.smtp.com -Port 22 -FromEmail support@company.com</code>
Related links	None

New-TransportConfiguration

Description	Helps you 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	<pre>New-TransportConfiguration [-AutoSelectTransport] <Boolean> [-RunLocal] <Boolean>[-UseAdvancedTransport] <Boolean> [-UseLanTransport] <Boolean> [[-UseSslforLan] <Boolean>] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<p>-----EXAMPLE 1-----</p> <pre>New-BackupFlag -RunLocal 1 -UseAdvancedTransport 1 UseLANTransport 0</pre> <p>This command returns a new transport configuration object with the values of RunLocal, UseAdvancedTransport.</p> <p>-----EXAMPLE 2-----</p> <pre>New-TransportConfiguration</pre> <p>This command returns a new transport configuration object with the configuration set to Automatic.</p>
Related links	Add-BackupJobTemplate Add-RestoreJobTemplate

New-VirtualAppliance

Description	Configures a VA for vRanger.
Syntax	<code>New-VirtualAppliance [-VAEntity] <Entity> [-RootPassword] <String> [[-IPAddressOverride] <String>] [[-ClusterDefault] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• VAEntity <Entity>: The VM Entity where the VA is deployed.• RootPassword <String>: Root Password for the VA.• IPAddressOverride <String>: New IP Address for the VA• ClusterDefault [<Boolean>]: Share the VA 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	<pre>\$va = Get-InventoryEntities -Type VirtualMachine where {\$_.Name -eq 'ABC-Vzva-1.1.1'} New-VirtualAppliance -VAEntity \$va -IPAddressOverride 10.6.50.19 -RootPassword vzroot1</pre> <p>This command adds a VA configuration to vRanger.</p>
Related links	Update-VirtualAppliance Remove-VirtualAppliance Get-VirtualApplianceConfig

New-WeeklySchedule

Description	Returns a Schedule object. The schedule provides a template weekly schedule that can be further customized before use.
Syntax	<code>New-WeeklySchedule [-ExecutionDays] <String> [-StartTime] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• ExecutionDays <String>: Days of the week on which the job runs: Monday, Tuesday, and so on. Separate multiple days with commas.• 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	<p>-----EXAMPLE 1-----</p> <pre>New-WeeklySchedule -ExecutionDays "Monday" -StartTime "2009-03-31 12:00:00AM"</pre> <p>This command returns weekly job schedule that runs on Mondays at midnight.</p> <p>-----EXAMPLE 2-----</p> <pre>New-WeeklySchedule -ExecutionDays "Monday,Tuesday" -StartTime "2009-03-31 12:00:00AM"</pre> <p>This command returns weekly job schedule that runs on Mondays and Tuesdays at midnight.</p>
Related links	None

New-YearlySchedule

Description	Returns a Schedule object. The schedule provides a template yearly schedule that can be further customized before use.
Syntax	<code>New-YearlySchedule [-ExecutionMonth] <Int> [-ExecutionDayOfMonth] <Int> [-StartTime] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• ExecutionMonth <Int>: Month in which the job runs—1 to 12.• ExecutionDayOfMonth <Int>: Day of month on which the job runs—1 to 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	<pre>New-YearlySchedule -ExecutionMonth 1 -ExecutionDayOfMonth 1 -StartTime "2009-03-31 12:00:00AM"</pre> <p>This command returns weekly job schedule that runs on the January 1 each year at midnight.</p>
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	<code>Remove-Allmount [<CommonParameters>]</code>
Parameters	<CommonParameters> : This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<pre>Remove-AllMount</pre> <p>This command removes all mounted paths.</p>
Related links	None

Remove-BackupGroupEntity

Description	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	<code>Remove-BackupGroupEntity [-BackupGroupId] <GUID> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>\$bg = Get-BackupGroupEntity where {\$_.Name -eq 'BG1'} Remove-BackupGroupEntity -BackupGroupId \$bg.Id</pre>
Related links	None

Remove-BackupGroupMember

Description	Removes members from a backup group.
Syntax	<pre>Remove-BackupGroupMember [-BackupGroupId] <GUID> [-MemberEntity] <Entity> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>\$vm = Get-InventoryEntities where {\$_.Name -eq 'myVM'} Remove-BackupGroupMember -BackupGroupId 5a62765e-6a04-4df4-af54-46e5a38ad876 -MemberEntity \$vm</pre>
Related links	None

Remove-Catalog

Description	Deletes the catalog data from the database for a list of savepoints.
Syntax	<pre>Remove-Catalog [-Savepoints] <GUID[]> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Remove-Catalog -Savepoints 5a62765e-6a04-4df4-af54-46e5a38ad876</pre> <p>Removes the catalog data associated with the savepoint from the database.</p>
Related links	Get-CatalogSearchData Get-CatalogStatus Set-Cataloging

Remove-DdbStorageUnit

Description	Removes the EMC® Data Domain® Boost (DD Boost™) repository from vRanger and deletes all savepoints in the repository from the disk. The storage unit is also removed from the DD Boost device.
Syntax	<code>Remove-DdbStorageUnit [-Repository] <Repository> [-Password] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• Repository <Repository>: The DD Boost 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	<pre>Remove-DdbStorageUnit -Repository \$repo -Password Password123</pre> <p>This command removes the repository from vRanger and deletes all associated savepoints from the disk. The storage unit is also removed from the DD Boost device.</p>
Related links	Remove-Repository

Remove-JobTemplate

Description	Removes a job template from vRanger.
Syntax	<code>Remove-JobTemplate [-Id] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>Remove-JobTemplate -Id \$guid</pre> <p>This command removes the job template with the specified ID from vRanger.</p>
Related links	Add-BackupGroupEntity Add-RestoreJobTemplate Get-JobTemplate

Remove-LinuxVolume

Description	This command is obsolete and has been replaced with Dismount-LinuxVolume .
Related links	Dismount-LinuxVolume

Remove-Repository

Description	Removes a repository from vRanger. Optionally, all savepoints in the repository can be deleted from the disk.
Syntax	<code>Remove-Repository [-Id] <String> [[-RemoveSavePoints] [<Boolean>]] [<i><CommonParameters></i>]</code>
Parameters	<ul style="list-style-type: none">• 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	<p>-----EXAMPLE 1-----</p> <pre>Remove-Repository -Id \$guid</pre> <p>This command removes the repository with the specified ID from vRanger. All savepoints associated with the repository remain un-touched.</p> <p>-----EXAMPLE 2-----</p> <pre>Remove-Repository -Id \$guid -RemoveSavePoints 1</pre> <p>This command removes the repository with the specified ID from vRanger. All savepoint files in the repository are removed from disk as well.</p>
Related links	Add-CIFSRepository Add-VirtualAppforLinuxFLR Add-HypervCluster Add-NVSDRepository Get-PhysicalMachineDiskMap

Remove-SavePoint

Description	Removes savepoints. If any of the specified savepoints has descendant savepoints, those savepoints are also removed.
Syntax	<code>Remove-SavePoint [-SavePointsToRemove] <SavePointList> [<i><CommonParameters></i>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>Remove-SavePoint -SavePointsToRemove \$splist</pre> <p>This command removes the list of savepoints specified.</p>
Related links	Get-RepositorySavePoint Remove-Repository

Remove-VirtualAppliance

Description	Removes the VA configuration from vRanger.
Syntax	<code>Remove-VirtualAppliance [-VAId] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• VAId <String>: The VA ID for the VA.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<pre>\$va = Get-VirtualApplianceConfig -VName KRN-Vzva-1.1.1 Remove-VirtualAppliance -VAId \$va.VirtualApplianceID</pre> <p>This command removed a VA configuration.</p>
Related links	New-VirtualAppliance Update-VirtualAppliance Get-VirtualApplianceConfig

Remove-VirtualApplianceConfiguration

Description	Removes the VA configuration from vRanger with an option to delete the VA from host.
Syntax	<code>Remove-VirtualApplianceConfiguration [-VAId] <String> [[-DeleteVM] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<pre>\$va = Get-VirtualApplianceConfig -VName KRN-Vzva-1.1.1 Remove-VirtualAppliance -VAId \$va.VirtualApplianceID -DeleteVM 1</pre> <p>This command removes a VA configuration from vRanger and also deletes the VM from Host.</p>
Related links	New-VirtualAppliance Update-VirtualAppliance Get-VirtualApplianceConfig

Run-JobsNow

Description	Runs jobs specified by their template IDs. Use the Get-JobTemplate cmdlet to view the template IDs of configured job templates.
Syntax	<code>Run-JobsNow [-TemplateIds] <String> [<CommonParameters>]</code>

Parameters	<ul style="list-style-type: none"> • TemplateIds <String>: The job template IDs. Separate multiple IDs with commas. • <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<pre>Run-JobsNow d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3</pre> <p>This command runs jobs specified by their templates.</p>
Related links	Add-BackupGroupEntity Add-RestoreJobTemplate Get-Job Get-JobTemplate

Run-ReplicationFailover

Description	Fails over the replication pass to the replica.
Syntax	<pre>Run-ReplicationFailover [-TemplateId] <String> [[-Sync] [<String>]] [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Run-ReplicationFailover -TemplateID d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3</pre> <p>This command runs replication failover based on the job template ID.</p>
Related links	Run-TestReplicationFailover Run-ResumeReplicationFailover

Run-ResumeReplicationFailover

Description	Resumes test replication failover based on the test replication task ID.
Syntax	<pre>Run-ResumeReplicationFailover [-TaskId] <String> [<CommonParameters>]</pre>
Parameters	<ul style="list-style-type: none"> • 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	<pre>Run-ResumeReplicationFailover -TaskID d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3</pre>
Related links	Run-TestReplicationFailover Run-ReplicationFailover

Run-TestReplicationFailover

Description	Runs test replication failover based on the replication job's template version ID.
Syntax	<code>Run-TestReplicationFailover [-TemplateVersionId] <GUID> [[-EnableNetwork] <Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<code>Run-TestReplicationFailover -TemplateVersionId d7ebc5c4-7916-44c4-a0c2-2412999735f3,d7ebc5c4-7916-44c4-a0c2-2412990005f3</code> This command runs test failover on a replication job.
Related links	Run-ResumeReplicationFailover Run-ReplicationFailover

Set-Cataloging

Description	Enables or disables global cataloging flag.
Syntax	<code>Set-Cataloging [-Catalog] <Boolean> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• 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	<code>Set-Cataloging -Catalog 1</code>
Related links	Remove-Catalog Get-CatalogSearchData Get-CatalogSearchData

Set-CBTonVM

Description	Enables or disables CBT for VMs if it is supported.
Syntax	<code>Set-CBTonVM [-VM] <VMEntity> [-CBT] [<Boolean>] [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none">• VM <VMEntity>: VM for which to enable CBT.• CBT [<Boolean>]: Enable or disable CBT.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"

Examples	<pre>\$vm = Get-InventoryEntities -type virtualmachine where{\$_.Name -eq 'xyz'}</pre> <pre>Set-CBTonVM -VM \$vm -CBT 1</pre> <p>This command enables CBT on the VM xyz if it is supported on the VMware® ESXi™ host.</p>
Related links	None

Set-LinuxVolume

Description	This command is 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	<pre>Set-MountPath [-Volume] <Volume> [-Path] <String> [<i><CommonParameters></i>]</pre>
Parameters	<ul style="list-style-type: none"> • Volume <Volume>: The volume to mount. • 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	<pre>Set-MountPath -volume \$v -path 'E:\'</pre> <p>This command mounts the specified volume in the savepoint to the E:\ drive.</p>
Related links	Get-PhysicalMachineDiskMap Get-RepositorySavePoint Get-RestoreDiskMap Get-VMDKVolume Remove-AllMount

Set-Resources

Description	Sets values for vRanger resources.
Syntax	<pre>Set-Resources [[-NumberTasksVRanger] <Int>] [[-NumberTasksLUN] <Int>] [[-NumberTasksHost] <Int>] [[-NumberTasksRepository] <Int>] [[-NumberTasksLocal] <Int>] [[-Timeout] <Int>] [[-MinimumSpaceHost] <Int>] [[-NumberTasksVA] <Int>] [<i><CommonParameters></i>]</pre>

Parameters	<ul style="list-style-type: none"> • NumberTasksVRanger [<Int>]: Maximum number of tasks running on vRanger (≥ 1). • NumberTasksLUN [<Int>]: Maximum number of tasks running off a LUN—1 to 5. • NumberTasksHost [<Int>]: Maximum number of tasks running on a host—1 to 4. • NumberTasksRepository [<Int>]: Maximum number of tasks running per Repository—≥ 1. • NumberTasksLocal [<Int>]: Maximum number of tasks running locally—1 to 20. • Timeout [<Int>]: Task Timeout in hours—1 to 999. • MinimumSpaceHost [<Int>]: Minimum space needed on host (GBs). • NumberTasksVA [<Int>]: Maximum number of tasks running on a VA—1 to 8. • <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<pre>Set-Resources -NumberTasksVRanger 10 -Timeout 2</pre> <p>Sets the maximum tasks run on vRanger and the Timeout for each task.</p>
Related links	Get-ConfigOption

Stop-vRangerJob

Description	Cancels a running job.
Syntax	<code>Stop-vRangerJob [-JobID] <String> [<CommonParameters>]</code>
Parameters	<ul style="list-style-type: none"> • 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	<p>-----EXAMPLE 1-----</p> <pre>Stop-vRangerJob -JobId 'F9C9E458-8E06-498D-8C2C-03CBE5515FD4'</pre> <p>This command cancels the job if it is running and returns the success or failure message.</p> <p>-----EXAMPLE 2-----</p> <pre>\$jobs = Get-Job where {\$_.JobState -eq 'Running'} Stop-vRangerJob \$jobs[0].Id</pre> <p>This command cancels the first running job in the list and returns the success or failure message.</p>
Related links	Remove-VirtualApplianceConfiguration Get-Job

Update-BackupJobTemplate

Description	Updates a backup job template. You also have the option of running the job immediately after the template is updated.
Syntax	<pre>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>]</pre>
Parameters	<ul style="list-style-type: none"> JobTemplate <BackupJobTemplate>: Existing job template. JobName [<String>]: Job name. JobDescription [<String>]: Job description. ExcludedVMList [<String>]: List of VMs 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 VM 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	<pre>get-jobtemplate backup where{\$_.jobname -eq 'my first backup job'} Update-BackupJobTemplate -JobName 'modified by vapi'</pre> <p>This command adds a backup job template. Because the RunJobNow flag is not specified, the job is not going to run immediately.</p>
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> [<i><CommonParameters></i>]
Parameters	<ul style="list-style-type: none">• TransportFailover <String>: Transport Failover order separated by commas to be attempted until the task is successful. Valid types are VA-HotAdd and VA-Network.• <CommonParameters>: This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, and -OutVariable. For more information, type: "get-help about_commonparameters"
Examples	<pre>Update-GlobalTransportFailover -TransportFailover VA-Network</pre> <p>This command updates the default Transport Failover order for replication jobs.</p>
Related links	Get-GlobalTransportFailover Get-TransportFailover

Update-HypervRestoreJobTemplate

Description	Updates a restore job template. You also have the option of running the job immediately after the template is updated.
Syntax	<pre>Update-HypervRestoreJobTemplate [-JobTemplate] <HypervRestoreJobTemplate> [[-JobName] <String>] [[-JobDescription] <String>] [[-SavePoint] <HypervSavePoint>] [-JobSchedule] <Schedule> [[-Flags] <NoFlags ForceOverwrite ForcePowerOn 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>]</pre>

Parameters

- **JobTemplate** *<HypervRestoreJobTemplate>*: Restore job template to update.
- **JobName** *<String>*: Job name.
- **JobDescription** *[<String>]*: Job description.
- **SavePoint** *<HypervSavePoint>*: 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>]*: Separate multiple email addresses with commas.
- **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 for the host is used.
- **SkipDisks** *[<List`1>]*: Sequence number—base 0—of the disks to skip from the savepoint. Separate multiple numbers with commas.
- **DiskLocations** *[<List`1>]*: Disk locations to use on the target host. Separate multiple values with commas. 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. Separate multiple values with 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. Separate multiple numbers with 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-----

```
$savepoint = Get-Repository |where{$_.name -eq  
'local'}|Get-RepositorySavepoint |where{$_.VmUuid -eq  
'12312312312313'}  
Get-JobTemplate restore |where{$_.jobname -eq 'my restore  
job'}|Update-HypervRestoreJobTemplate -savepoint $savepoint
```

This command updates a restore job template to restore a different savepoint.

Related links

[Add-RestoreJobTemplate](#)
[Get-JobTemplate](#)
[Get-Job](#)

Update-Inventory

Description	Refreshes the inventory.
Syntax	Update-Inventory [<i><CommonParameters></i>]
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. You also have the option of running the job immediately after the template is updated.
Syntax	Update-ReplicationJobTemplate [-JobTemplate] <i><ReplicationJobTemplate></i> [[-JobName] <i><String></i>] [[-JobDescription] <i><String></i>] [[-TargetHost] <i><HostEntity></i>][[-ReplicateName] <i><String></i>] [[-Type] <i><String></i>] [[-TargetNetworks] <i><String></i>] [[-PrimaryDatastore] <i><String></i>] [[-TargetDatastores] <i><String></i>] [[-JobSchedule] <i><Schedule></i>] [[-TransportFailover] <i><String></i>] [[-NotificationList] <i><String></i>] [[-Flags] <i><NoFlags ReplicatePoweredOnOnly CheckDestinationFreeSpace UseCompression UpdateNotes EnableGuestQuiescing EnableABM AllFlags></i>] [[-RunJobNow] <i><Boolean></i>] [[-IncludedDisk] <i><String></i>] [<i><CommonParameters></i>]

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>*]: Name of replicated VM.
- **Type** [*<String>*]: Replication type. Valid types are Differential and Hybrid.
- **TargetNetworks** *<String>*: Networks to use on the target host. Separate multiple values with commas.
- **PrimaryDatastore** [*<String>*]: Primary Datastore on target host on which the replicated VM is located.
- **TargetDatastores** [*<String>*]: Datastores to use on the target host. Separate multiple values with commas.
- **JobSchedule** [*<Schedule>*]: The recurrence schedule on which the job runs.
- **TransportFailover** [*<String>*]: Transport Failover order separated by commas to be attempted until the task is successful. Valid types are VA-HotAdd and VA-Network.
- **NotificationList** [*<String>*]: Email notification list for the job. Separate multiple email addresses with commas.
- **Flags** [*<ReplicationFlags>*]: Replication flags: NoFlags, ReplicatePoweredOnOnly, CheckDestinationFreeSpace, UseCompression, UpdateNotes, EnableGuestQuiescing, AllFlags. Separate multiple flags with pipes (|).
- **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. You also have the option of running the job immediately after the template is updated.
-------------	--

Syntax	<pre>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>]</pre>
--------	--

Parameters

- **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. Separate multiple flags with pipes (|).
- **NotificationList** [<String>]: Email notification list for the job. Separate multiple email addresses with commas.
- **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. Separate multiple numbers with commas.
- **TargetDatastores** [<String>]: Datastores to use on the target host. Separate multiple values with commas. 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. Separate multiple values with 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; 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. Separate multiple numbers with commas.
- **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

This command is obsolete. Use [Update-VirtualApplianceConfiguration](#) instead.

Related links

[New-VirtualAppliance](#)
[Remove-VirtualAppliance](#)
[Get-VirtualApplianceConfig](#)

Update-VirtualApplianceConfiguration

Description

Updates the VA 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 VA instead of the one retrieved from the VMware® vCenter™ API.
- **ClusterDefault** **[<Boolean>]**: If the selected host belongs to a cluster, share the VA among all hosts in the cluster.
- **CPUAllocation** **[<Int>]**: Allocate CPU.
- **MemorySizeInMB** **[<Int>]**: Assign memory in MB.
- **ScratchDiskDatastore** **[<Datastore>]**: Select a datastore for the scratch disk.
- **ScratchDiskSizeInKB** **[<Int>]**: Assign scratch disk in KB.
- **NewPassword** **[<String>]**: New password for the VA, if this word is not specified, the password remains unchanged.
- **UseVAForLFLR** **[<Boolean>]**: Use VA for Linux® FLR.
- **DeleteScratchDisk** **[<Boolean>]**: If set to **true**, deletes the scratch disk—the default is **false**.
- **VAreconfigurationStatus** **[<Boolean>]**: Retrieves VA deployment status if set to true. If the parameter is not passed in, the default is false.
- **<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 -VName KRN-Vzva-1.1.1 Update-  
VirtualApplianceConfiguration -VirtualApplianceEntity $va  
-CPUAllocation 2 -VAreconfigurationStatus 1
```

This command updates the VA configuration and sets CPUAllocation to 2. Setting VAreconfigurationStatus to true retrieves the status of the VA reconfiguration in progress. Default is false.

Related links

[New-VirtualAppliance](#)
[Remove-VirtualAppliance](#)
[Get-VirtualApplianceConfig](#)

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit <https://www.quest.com/company/contact-us.aspx> or call +1-949-754-8000.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.

Writer's note: If the product does not have support through normal channels, apply the DoNotPublish condition to the preceding text and delete it from the next paragraph. Modify the statement to reflect the support options available for your release. If no support is available through any channel, use only the first sentence.

This version of vRanger does not qualify for product support by Quest Support. Visit the [support portal](#) for documentation and Knowledge Base articles on previous versions, visit the product community, consult the product documentation, or [other method] for support.