

VMware Data Recovery

Administrator's Guide

EN-000193-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

© 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware, the VMware "boxes" logo and design, Virtual SMP, and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book	5
1 Understanding VMware Data Recovery	7
Backing Up Virtual Machines	7
Volume Shadow Copy Service	8
Deduplication Store Benefits	9
2 Installing VMware Data Recovery	11
VMware Data Recovery System Requirements	11
Install the Client Plug-in	12
Install the Backup Appliance	12
Add a Hard Disk to the Backup Appliance	13
3 Configuring VMware Data Recovery	15
Power On the Backup Appliance	15
Configure the Backup Appliance	16
Use the Getting Started Wizard	16
Creating a Backup Job	17
Restoring Virtual Machines	19
Troubleshooting VMware Data Recovery	21
Index	23

About This Book

The *VMware Data Recovery Administrator's Guide* contains information about establishing backup solutions for small and medium businesses.

Intended Audience

This book is for anyone who wants to provide backup solutions using VMware Data Recovery. The information in this book is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to docfeedback@vmware.com.

Technical Support and Education Resources

The following technical support resources are available to you. To access the current version of this book and other books, go to <http://www.vmware.com/support/pubs>.

Online and Telephone Support

To use online support to submit technical support requests, view your product and contract information, and register your products, go to <http://www.vmware.com/support>.

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to http://www.vmware.com/support/phone_support.html.

Support Offerings

To find out how VMware support offerings can help meet your business needs, go to <http://www.vmware.com/support/services>.

VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to <http://www.vmware.com/services>.

Understanding VMware Data Recovery

1

VMware® Data Recovery creates backups of virtual machines without interrupting their use or the data and services they provide. Data Recovery manages existing backups, removing backups as they become older. It also supports deduplication to remove redundant data.

Data Recovery is built on the VMware vStorage API for Data Protection. It is integrated with VMware vCenter Server, allowing you to centralize the scheduling of backup jobs. Integration with vCenter Server also enables virtual machines to be backed up, even when they are moved using VMware VMotion™ or VMware Distributed Resource Scheduler (DRS).

Data Recovery uses a virtual machine appliance and a client plug-in to manage and restore backups. The backup appliance is provided in open virtualization format (OVF). The Data Recovery plug-in requires the VMware vSphere Client.

Backups can be stored on any virtual disk supported by VMware ESX™. You can use storage area networks (SANs), network attached storage (NAS) devices, or Common Internet File System (CIFS) based storage such as SAMBA. All backed-up virtual machines are stored in a deduplicated store.

VMware Data Recovery supports the Volume Shadow Copy Service (VSS), which provides the backup infrastructure for certain Windows operating systems.

This chapter includes the following topics:

- [“Backing Up Virtual Machines,”](#) on page 7
- [“Volume Shadow Copy Service,”](#) on page 8
- [“Deduplication Store Benefits,”](#) on page 9

Backing Up Virtual Machines

During a backup, Data Recovery creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

Data Recovery can concurrently back up a maximum of eight virtual machines. To start multiple backups, CPU utilization must be less than 90 percent. Due to memory constraints, Data Recover does not support using more than two backup destinations simultaneously. If more than two backup destinations must be used, configure them to be used at different times.

For virtual machines created in vSphere 4.0, the Data Recovery appliance creates a quiesced snapshot of the virtual machine during the backup. The backups use the changed block tracking functionality on the ESX hosts. For each virtual disk being backed up, it checks for a prior backup of the virtual disk. It uses the change-tracking functionality on ESX hosts to obtain the changes since the last backup. The deduplicated store creates a virtual full backup based on the last backup image and applies the changes to it.

NOTE These optimizations do not apply to virtual machines created with VMware products prior to vSphere 4.0. For example, change tokens are not used with virtual machines created with Virtual Infrastructure 3.5 or earlier. As a result, virtual machines created with earlier VMware versions take longer to back up.

If duplicate parts of a virtual machine are found, a record of the information is stored rather than storing the information twice. Deduplication can provide significant space savings. Operating system files are often identical among virtual machines running the same operating system. To maximize deduplication, back up similar virtual machines to the same destination. The virtual machines do not need to be backed up during the same job.

Data Recovery uses the vSphere licensing infrastructure to ensure that all virtual machines that are protected by Data Recovery have appropriate licensing. Valid vSphere licensing includes Essential Plus, Advanced, Enterprise, or Enterprise Plus licenses.

Each Data recovery backup appliance can protect a total of 100 virtual machines. It is possible to create backup jobs that are configured to protect more than 100 virtual machines, but the backup appliance only protects 100 virtual machines and any additional virtual machines are omitted. It is possible to protect more than 100 virtual machines by installing additional backup appliances, but different backup appliances do not share information about backup jobs. As a result, it is possible to establish unintended configurations. For example, two Data Recovery backup appliances could be configured to protect a folder containing 200 virtual machines, but it is likely that some of the virtual machines would be backed up twice and some would not be backed up at all.

Volume Shadow Copy Service

VMware Data Recovery uses the Microsoft Windows Volume Shadow Copy Service (VSS), which provides the backup infrastructure for certain Windows operating systems, as well as a mechanism for creating consistent point-in-time copies of data known as shadow copies.

VSS produces consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware. VSS support is provided with VMware Tools, which runs in the guest operating system. VMware provides a VSS Requestor and a VSS Snapshot Provider (VSP). The Requestor component is available inside a supported guest and responds to events from an external backup application. The Requestor also controls the progress of backup operations inside the guest and interacts with the VSP. The Requestor is instantiated by the VMware Tools service when a backup process is initialized. The VSP is registered as a Windows service and notifies Data Recovery of provider-specific events during a VSS backup.

VSS is supported on virtual machines with the following guest operating systems:

- Windows Server 2003, 32 bit and 64 bit
- Windows Vista, 32 bit and 64 bit
- Windows Server 2008, 32 bit and 64 bit

For virtual machines with Windows operating systems that do not support VSS, VMware Tools uses the default LGTO SYNC driver. For other guest operating systems, VMware Tools uses crash-consistent quiescing.

Several features in the Windows Server 2003 operating systems use the Volume Shadow Copy Service, including Shadow Copies for Shared Folders and Backup. Shadow copies can be used for a number of purposes such as:

- Creating consistent backups of open files and applications
- Creating shadow copies for shared folders

- Quickly recovering and restoring files and data
- Creating transportable shadow copies using a hardware provider — for backup, testing and data mining scenarios

Because Data Recovery uses VSS, Data Recovery can create snapshots while ensuring application consistency. This means that applications write to disk any important data that is currently in memory, making sure that a later restore of that virtual machine can restore the application back into a consistent state.

For more information about which Windows virtual machines use the Volume Shadow Copy Service, see the Virtual Machine Backup Guide. Detailed information about VSS can be found at <http://technet.microsoft.com/en-us/library/cc785914.aspx>.

Deduplication Store Benefits

The deduplication store technology used by VMware Data Recovery was developed by VMware and provides tight integration. The deduplication technology evaluates patterns to be saved to restore points and checks to see if identical sections have already been saved.

Because VMware supports storing the results of multiple backup jobs to use the same deduplication store, to maximize deduplication rates, ensure that similar virtual machines are backed up to the same destination. While backing up similar virtual machines to the same deduplication store may produce increased space-savings, the similar virtual machines do not need to be backed up during the same job. Deduplication is evaluated for all virtual machines stored, even if some are not currently being backed up.

Data Recovery is designed to support deduplication stores that are up to one terabyte in size and each backup appliance is designed to support the use of two deduplication stores. Data Recovery does not impose limits on the size of deduplication stores or number of deduplication stores, but if more than two stores are used or as the size of a store exceeds one terabyte, performance may be affected.

Installing VMware Data Recovery

VMware Data Recovery uses a plug-in to the vSphere Client and a backup appliance to store backups to destinations such as hard disks.

Before you can begin using Data Recovery, you must complete the installation process, beginning with ensuring that your environment includes resources that meet the Data Recovery system requirements.

This chapter includes the following topics:

- [“VMware Data Recovery System Requirements,”](#) on page 11
- [“Install the Client Plug-in,”](#) on page 12
- [“Install the Backup Appliance,”](#) on page 12
- [“Add a Hard Disk to the Backup Appliance,”](#) on page 13

VMware Data Recovery System Requirements

Before installing VMware Data Recovery, ensure the system and storage requirements are available in your environment.

Data Recovery requires vCenter Server and the vSphere Client. Data Recovery does not work with similar VMware products such as VirtualCenter Server. You can download the vSphere Client from your vCenter Server. Virtual machines must be running on an ESX 4.0 server or later.

You can store backups on any virtual disk supported by ESX. You can use technologies such as storage area networks (SANs) and network attached storage (NAS) devices. Data Recovery also supports Common Internet File System (CIFS) based storage such as SAMBA.

The amount of storage required varies, depending on how much deduplication can save disk space as a result of running similar virtual machines, but Data Recovery requires a minimum of 10 GB of space as a minimum. The more diverse the set of virtual machines to be protected, the more space is required for each virtual machine. The amount of space required is also affected by the frequency of backup, the length of time the backups are kept, and the number of virtual machines to be backed up.

See the most recent vSphere documentation for information about setting up a vSphere 4.0 environment including ESX Servers, vCenter Server, and the vSphere client.

Install the Client Plug-in

You must install the client plug-in before you can manage VMware Data Recovery.

Prerequisites

Before you can install the Data Recovery plug-in, you must have vCenter Server running in your environment, and you must install the vSphere Client, which you can download from any vCenter Server. The Data Recovery plug-in connects to the backup appliance using port 22024. If there is a firewall between the client and the backup appliance, port 22024 must be open before Data Recovery can be managed with the vSphere Client.

Procedure

- 1 Run the plug-in installer `VMwareDataRecoveryPlugin.msi`.
- 2 Follow the prompts of the installation wizard.
- 3 Start the vSphere Client, and log in to a vCenter Server.
- 4 Select **Plugins > Manage Plugins** and make sure that the Data Recovery plug-in is enabled.

You can now use the client plug-in to manage Data Recovery. If the Data Recovery is not registered in the vSphere Client, restart the client.

What to do next

You may now want to complete the task [“Install the Backup Appliance,”](#) on page 12.

Install the Backup Appliance

You must install the backup appliance so Data Recovery can complete backup tasks. You use the vSphere Client to deploy the backup appliance.

Prerequisites

To install the backup appliance, you must have a vCenter Server and an ESX server running in your environment.

Procedure

- 1 From the vSphere Client, select **File > Deploy OVF Template**.
- 2 Select **Deploy from file**, and then browse to `VMwareDataRecovery.ovf` and select it.
- 3 Review the OVF file details.
- 4 Review the End User License Agreement. If you agree to the terms, accept them.
If you do not accept the terms, you cannot complete the process.
- 5 Select a location for the backup appliance in the vSphere inventory.
You can optionally rename the backup appliance.
- 6 Select the host or cluster to which the backup appliance is to be deployed.
- 7 Select a datastore to store the virtual machine files.
- 8 Review the deployment settings and click **Finish**.

The backup appliance is now deployed into your environment.

What to do next

You can save backups on network storage or on hard disks. If you are going to store backups on a hard disk, you may now want to complete the task [“Add a Hard Disk to the Backup Appliance,”](#) on page 13. Otherwise you may now want to learn about [Chapter 3, “Configuring VMware Data Recovery,”](#) on page 15.

Add a Hard Disk to the Backup Appliance

You can store backups to a hard disk that has been added to the backup appliance. Hard disks provide faster backup performance compared to other destinations such as CIFS shares.

Prerequisites

If you are adding a hard disk, you must have installed the backup appliance and the Data Recovery plug-in for the vSphere Client.

Procedure

- 1 Start the vSphere Client and log in to the vCenter Server that manages the backup appliance.
- 2 Select **Inventory > VMs and Templates**.
- 3 In the inventory, right-click the backup appliance virtual machine and select **Edit Settings**.
- 4 In the Hardware tab, click **Add**.
- 5 Select **Hard Disk** and click **Next**.
- 6 Select **Create a new virtual disk** and click **Next**.
- 7 Specify the disk size and other options and click **Next**.
- 8 Specify the advanced options for the virtual disk and click **Next**.
- 9 Click **Finish** to create the disk.

The disk is now added to the backup appliance and can be used as a destination for backups. If the backup appliance is powered on when the hard disk is added, the hard disk is not recognized until the backup appliance is rebooted. Therefore, if the backup appliance is powered on, reboot it to complete the addition of the hard disk.

What to do next

You may now want to learn about [Chapter 3, “Configuring VMware Data Recovery,”](#) on page 15.

Configuring VMware Data Recovery

To configure Data Recovery, you connect the backup appliance to vCenter Server and specify backup configurations.

Common tasks involved with establishing backup configurations include:

- Powering on the backup appliance.
- Connecting the backup appliance to the vCenter Server.
- Configuring Data Recovery.
- Establishing backup jobs, including required resources, which may include adding network shares or formatting volumes.

This chapter includes the following topics:

- [“Power On the Backup Appliance,”](#) on page 15
- [“Configure the Backup Appliance,”](#) on page 16
- [“Use the Getting Started Wizard,”](#) on page 16
- [“Creating a Backup Job,”](#) on page 17
- [“Restoring Virtual Machines,”](#) on page 19
- [“Troubleshooting VMware Data Recovery,”](#) on page 21

Power On the Backup Appliance

You must power on the virtual machine backup appliance to perform backups.

Prerequisites

Before powering on the backup appliance, you must have completed the process of [“Install the Backup Appliance,”](#) on page 12.

Procedure

- 1 In the vSphere Client, select **Inventory > VMs and Templates**.
- 2 In the inventory, right-click the virtual machine to use as the backup appliance and select **Power On**.
- 3 After the virtual machine is powered on, right-click the backup appliance virtual machine and choose **Open Console**.

The console window for the backup appliance appears.

- 4 Provide the username and credentials for this system.

If this is the first time logging on to the backup appliance, the default credentials are username: root, password: vmw@re.

- 5 If the root account password has not been changed from the default, use the `passwd` command to change the password for the root account to a strong password of your choosing.
- 6 Close the console window.

The backup appliance is left powered on, ready to complete backup tasks.

Configure the Backup Appliance

You can configure networking settings or reboot the backup appliance, as required, using the web interface.

Prerequisites

Before you can configure the backup appliance, it must be powered on.

Procedure

- 1 Enter the URL for the backup appliance in a web browser.
The URL for the backup appliance is displayed on the appliance console. To view the appliance console, open it from the vSphere Client.
- 2 Provide the username and password for the administrator.
- 3 Click the **System** tab to gather information about the appliance or click **Reboot** or **Shutdown**, as required.
- 4 Click the **Network** tab and click **Status** for information about current network settings.
- 5 Click the **Network** tab and click **Address** to configure network settings. You can configure the backup appliance to obtain its address from a DHCP or you can manually configure IP settings.
- 6 Click the **Network** tab and click **Proxy** to configure proxy settings. You can configure the backup appliance to use a proxy server and provide the proxy server's name or IP address and port.

Use the Getting Started Wizard

Use the getting started wizard to establish an initial system configuration that is used to begin backing up virtual machines to restore points.

Prerequisites

Before using the Getting Started Wizard, you must complete the process described under [Connect the Backup Appliance to vCenter Server](#).

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 In the Credentials page, enter a username and password and click **Next**.
Data Recovery uses this information to connect to vCenter to perform backups, so the specified user account must have administrative rights.
- 3 In the Backup Destinations page, select a backup destination from the list of choices.

- 4 In the Backup Destinations page, select the tasks that you want to perform.
 - To rescan the SCSI bus for new SCSI devices, click **Refresh**.
 - To format a virtual disk that has been added to the appliance, click **Format**. After formatting completes, the disk appears as `scsi:x:y`. For disks that already contain data, use **Mount** rather than format.
 - To mount a disk containing an existing deduplication store, click **Mount**.
 - To mount the CIFS share, click **Add Network Share** and provide credentials. These credentials are stored in the appliance, so remounting is completed automatically if the appliance is rebooted.
- 5 Click **Next**.

The initial system configuration is now complete and the Create a New Backup Job wizard opens by default. Use the Create a New Backup Job wizard, as described in “[Creating a Backup Job](#),” on page 17 to create a backup job.

Creating a Backup Job

You can create backup jobs that include which virtual machines to backup, where to store the backups, and for how.

Virtual Machines

You can specify collections of virtual machines, such as all virtual machines in a datacenter, or select individual virtual machines. If an entire resource pool, host, datacenter, or folder is selected, any new virtual machines in that container are included in subsequent backups. If a virtual machine is selected, any disk added to the virtual machine is included in the backup. If a virtual machine is moved from the selected container to another container that is not selected, it is no longer part of the backup.

Destination

You can store backups on network shares, in VMDKs, or on RDMs. If you are storing backups on a network share and the network share on which you want to store the backup is not available, you can add a network share. For more information, see “[Add a Network Share](#),” on page 18. You must format and partition VMDKs and RDMs to store backups. You can format or partition destinations that are not yet formatted or partitioned. For more information, see “[Formatting a Volume](#),” on page 19.

Backup Window

By default, backup jobs run at night on Monday through Friday and at any time on Saturday and Sunday. Data Recovery attempts to back up each virtual machine in a job once a day during its backup window. If the backup timeframe for the backup window passes while the backup is in progress, the backup is stopped. The backup will be restart when the backup window opens. This means that if there are too many virtual machines for Data Recovery to back them all up during the first specified window, some virtual machines may not be backed up. Eventually Data Recovery will complete backup of all virtual machines and subsequent backups typically fit within one backup window. If some machines are not backed up during a window, those machines are given higher priority during subsequent backup windows. This helps ensure that all virtual machines are backed up as often as the back windows and resources allow, and prevents the case where some virtual machines are always backed up and some are never backed up.

Retention Policy

Data Recovery backups are preserved for a variable period of time. You can choose to keep more or fewer backups for a longer or shorter period of time. Keeping more backups consumes more disk space, but also provides more points in time to which you can restore virtual machines. As backups age, some are automatically deleted to make room for new backups. You can use a predefined retention policy or create a custom policy. The backup policy is once a day during the backup window.

Ready to Complete

Review the settings for the backup job. This page includes information including:

- Which virtual machines will be backed up by this job.
- Where the backups for the specified virtual machines will be stored.
- The schedule on which virtual machines will be backed up.
- The number of backups that will be kept for the segments of time. For example, the number of backups that will be kept for each month.

Use the Backup Job Wizard

Use the Backup Job Wizard to specify which virtual machines are to be backed up and when this can occur.

Prerequisites

Before using the Backup Job Wizard, you must establish a VMware Data Recovery configuration. This can be completed using the Getting Started Wizard, as described under [“Use the Getting Started Wizard,”](#) on page 16.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 Click **New Backup Job...** to launch the Backup Job wizard.
- 3 In the Virtual Machines page, select individual virtual machines or containers that contain virtual machines to be backed up and click **Next**.
- 4 In the Backup Window page, accept the default times or specify alternate backup windows and click **Next**.
- 5 In the Retention Policy page, accept the default retention policy or specify an alternate retention policy and click **Next**.
- 6 In the Ready to Complete page, reviewed the summary information for the backup job and click **Next**.

Add a Network Share

You can establish a network share on which backups are stored.

Provide information about a network share on which VMware Data Recovery can store backups. Information typically required includes:

- URL - The location of the network share. This must be specified with the IP address. Server names are not supported.
- User name - The user name for an account with the required write privileges for the network share.
- Password - The password for the user account.

Formatting a Volume

VMware Data Recovery can store backups on network volumes, VMDKs, and RDMs. Networked volumes might not require formatting, but VMDKs and RDMs must be formatted before they can be used.

Formatting a volume automatically formats and partitions the space. As a result, any data that is stored in this space is erased. As required, format the volume you intend to use for backup storage.

Restoring Virtual Machines

You can specify which virtual machines to restore and how they are restored using the Virtual Machine Restore wizard.

Data Recovery provides means to test how a virtual machine would be restored and to actually carry out restore operations. Restore rehearsals create virtual machines from restore points. Virtual machines from restore rehearsals do not replace current virtual machines, but virtual machines create through restore rehearsals do provide a way to ensure that virtual machine backups are being created as expected and that they can be successfully restored. Actually restoring virtual machines returns specified virtual machines to a selected previous state.

For both restore rehearsals and restores, the Restore Virtual Machines wizard provides pages that allow you to configure from where and to where virtual machines are restored.

Source Selection

When choosing a source, select from the tree view of backed up vSphere objects. Select those virtual machines and virtual disks to be restored. You can use filters to view a subset of all available choices. Much like with creating back up jobs, you can specify collections of virtual machines, such as all virtual machines in a datacenter, or select individual virtual machines or vmdk files to restore. If multiple restore points are selected for a single virtual machine, Data Recovery restores the virtual machine to the most recent restore point selected.

Destination Selection

This page provides a tree view of the location to which backed up vSphere objects will be restored and how those objects will be configured when they are restored. If your inventory hierarchy changed since the time of the backup, inventory object paths that no longer exist are shown as grayed out. You must move virtual machine files that were backed up from locations that no longer exist to valid destinations before you can perform the restore operation. You can reconfigure options such as:

- The datastore and virtual disk node to which the files will be restored.
- Whether the configuration will be restored. If configuration is not restored, configuring some other options may not be supported. For example, if the configuration is not restored, it may be possible to configure whether the virtual machine will be powered on, but not whether the NIC will be connected.
- Whether the NIC will be connected.
- Whether the virtual machine will be powered on.

It is possible to move virtual machines and VMDKs to different locations either by dragging and dropping them, or by selecting new destinations from the popup tree. To see more information about the existing inventory, click the link at the top of the page.

To clone a virtual machine, rename the virtual machine you are restoring.

If the default credentials provided for backup do not have privileges for restore, you can specify alternate credentials.

Ready to Complete

Review the settings for the backup job. This page includes a tree-style representation of what will be restored and summary information. The tree-style representation includes information such as:

- Object names.
- When the restore point was created.
- Which datastore will be used as the destination for restored virtual machines or virtual disks.
- Virtual disk node information.
- Whether the configuration will be restored.
- Whether the NIC will be connected.
- Whether the virtual machine will be powered on.

The summary contains information such as:

- How many virtual machines will be overwritten.
- How many virtual machines will be created.
- How many virtual disks will be overwritten.
- How many virtual disks will be created.
- The total amount of data that will be restored.

NOTE If there is insufficient space on the destination datastore to complete the restore, a warning is displayed. Specify alternate datastores with increased capacity or accept the possibility that restores may not complete as expected.

Complete a Restore Rehearsal

Complete a restore rehearsal to confirm that a virtual machine is being backed up as expected and that a successful restore operation would complete as expected.

Prerequisites

Before you can complete a restore rehearsal, you must have configured VMware Data Recovery and have at least one backup.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**
- 2 Right-click a virtual machine that has a backup and select **Restore Rehearsal**.
The Virtual Machine Restore Wizard appears displaying the ready to complete page.
- 3 Click **Restore** to complete the restore rehearsal or click **Back** to modify settings.

A version of the virtual machine is restored to the inventory.

What to do next

Next you may want to delete the virtual machine that was created in testing the restore process.

Restore Virtual Machines from Backup

Restore virtual machines to a previous backup state using the Restore Virtual Machines wizard.

Prerequisites

Before you can restore virtual machines, you must have configured VMware Data Recovery and have at least one backup to which to restore.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 Enter the virtual machine name or IP address of the backup appliance and click **Connect**.
- 3 Click the **Restore** tab and click the **Restore** link to launch Virtual Machine Restore Wizard.
The Restore Virtual Machines wizard appears.
- 4 On the Source Selection page, specify a source from which to restore virtual machines and click **Next**.
- 5 On the Destination Selection page, specify how restored machines will be configured and click **Next**.
- 6 On the Ready to Complete page, review the configuration and click **Finish**.

The virtual machines are restored as specified in the wizard.

Troubleshooting VMware Data Recovery

If you have connection or configurations problems with Data Recovery, you can try to resolve them using by troubleshooting.

Problem	Possible Solution
Unable to connect to the backup virtual machine.	Check that the virtual machine is listed under the Host and Clusters view of the inventory. Make sure that the IPv4 address of the Data Recovery appliance is correct.
The NFS share is not working as expected.	NFS is only supported if the share is presented by an ESX Server and the VMDK is assigned to the appliance. NFS shares cannot be mapped directly to the appliance.
Data Recovery has crashed. Is the system state fine?	Because the state of the appliance is stored in the deduplicated store, it can be restored. Reinstall the Data Recovery appliance to the ESX host, and configure the appliance to point to the existing deduplicated store.
The backup appliance is connected to vCenter Server and a crash has occurred.	If the vSphere Client crashes after applying changes, restart the vSphere Client and reconnect to the backup appliance.

A valid network name is entered, but Data Recovery does not connect.

In some cases, name resolution might not work. Try using the IP address for the desired target.

Backup and restore operations are not completing as expected.

An integrity check may have discovered a problem with the integrity of the deduplication store.

The integrity of new backups is checked each day, and the entire deduplication store is checked once a week. If problems are found during the integrity check, the deduplication store is locked. As a result, no backups or restores can be performed until the issues reported by the integrity check are fixed. To resolve this issue, select the problematic restore points on the restore tab, and click Mark for Delete. These restore points are deleted during the next integrity check, after which the deduplication store is unlocked.

If no integrity check problem has been identified, the issue may be due to an excess of jobs. Data Recovery limits the number of jobs that can run to help prevent systems from becoming overloaded and failing to make progress. Some of the limits include:

- Maximum of eight backup jobs can run at the same time.
- Maximum of eight restore jobs can run at the same time.
- Processor must not exceed 90 percent utilization if a job is to start.
- Deduplication store must have at least 5GB of storage space available for each job.

If any of these limits are exceeded, new jobs do not start.

If you have problems that cannot be resolved using these troubleshooting tips, you can open a service request with VMware technical support. Before contacting technical support, consider gathering Data Recovery log files and hidden logs and executing the log gathering script.

Index

A

- adding
 - network share **18**
 - storage **13**

B

- backup
 - process **7**
 - scaling **7**
- backup appliance
 - configuring **16**
 - installing **12**
 - power on **15**
- backup job
 - creating **17, 18**
 - options **17**
- backup job wizard, using **18**

C

- client, installing **12**
- configuring
 - backup appliance **16**
 - data recovery **15**
- creating, backup job **17**

D

- data recovery
 - configuring **15**
 - prerequisites **11**
 - scaling **11**
- deduplication
 - best practices **9**
 - scaling **9**

F

- firewalls **12**

G

- getting started wizard, using **16**

I

- installing
 - backup appliance **12**
 - client **12**
 - data recovery **11**

- introducing, data recovery **7**

L

- licensing **7**

N

- network share, adding **18**

R

- restore rehearsal **19, 20**
- restoring, virtual machines **19, 21**

S

- scaling
 - backup **7**
 - data recovery **11**
 - deduplication **9**
 - storage, adding **13**
 - supported storage **7**

T

- troubleshooting **21**

U

- using, getting started wizard **16**

V

- virtual machines, restoring **19, 21**
- Volume Shadow Copy Service, *See also* VSS
- volumes, formatting **19**
- VSS
 - benefits **8**
 - support **8**
 - understanding **8**

