VMware Data Recovery Administration Guide

Data Recovery 1.1

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

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

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

Contents

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

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 Requester 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
- Windows Server 2008 R2
- Windows 7, 32 bit and 64 bit

Data Recovery uses different quiescing mechanisms depending on the guest operating system that you run in your virtual machines.

Table 1-1. Quiescing Mechanisms Used by Data Recovery

Guest Operating System	Driver and Quiescing Type Used
Windows XP 32-bit Windows 2000 32-bit	Sync Driver File-system consistent quiescing
Windows 2003 32-bit/64-bit	VMware VSS component Application-consistent quiescing
Windows Vista 32-bit/64-bit Windows 7 32-bit/64-bit Windows 2008 32-bit/64-bit Windows 2008 R2	VMware VSS component File-system consistent quiescing
Other guest operating systems	Crash-consistent quiescing

In most cases, the quiescing mechanisms provided with Data Recovery will properly quiesce applications. If your environment includes applications or operating systems that do not respond to included quiescing mechanisms as expected, Data Recovery supports the use of custom quiescing scripts. Deploy and run the custom quiescing scripts inside the protected virtual machine.

Table 1-2. Locations of Custom Quiescing Scripts

Guest Operating System	Script	Location of Script on Virtual Machine
Windows	Pre-freeze	`C:\Program Files\VMware\VMware Tools\backupScripts.d` All scripts are invoked in ascending alphabetical order with freeze as the first argument.
	Post-thaw	`C:\Program Files\VMware\VMware Tools\backupScripts.d` All scripts are invoked in descending alphabetical order with thaw or freezeFail as the first argument.
Other	Pre-freeze	`/usr/sbin/pre-freeze-script`
	Post-thaw	`/usr/sbin/post-thaw-script`

When running the scripts, you can also use the SYNC driver or VSS components on those virtual machines that support them.

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.

There are several processes that the deduplication store completes including integrity check, recatalog, and reclaim.

Integrity check

This operation is performed to verify and maintain data integrity on the deduplication store. This operation runs automatically on a daily basis, though the complete integrity check is performed once a week. In addition, the integrity check can be executed manually. Normally, the backup and restore operations are allowed from the deduplication store while the integrity check is in progress. If a restore point is manually marked for delete, backups are not allowed during integrity check but restore operations are allowed. If damaged restore points are found in the deduplication store during integrity check, a manual integrity check must be run after marking the damaged restore points for delete. During this manually run integrity check, backups and restore are not allowed.

Recatalog

This operation is performed to ensure that the catalog of restore points is synchronized with the contents of the deduplication store. This operation runs automatically when there is an inconsistency detected between the catalog and the deduplication store. While the recatalog operation is in progress, no other operation is allowed on the deduplication store.

Reclaim

This operation is performed to reclaim space on the deduplication store. This can be a result of the Data Recovery appliance enforcing the retention policy and deleting expired restore points. This operation runs automatically on a daily basis or when a backup job requires more space than is available on the deduplication store. While the reclaim operation is in progress, backups to the deduplication store are not allowed, but restore operations from the deduplication store are allowed.

Installing VMware Data Recovery 2

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.

Data Recovery is composed of a set of components that run on different machines.

- The client plug-in is installed on a computer that will be used to manage Data Recovery.
- The backup appliance is installed on an ESX host.
- The optional File Level Restore (FLR) client is installed in a windows virtual machine. For more information on FLR, see "Understanding File Level Restore," on page 24.

This chapter includes the following topics:

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 to be backed up and the backup appliance must both be running on ESX 4 or later or ESX 4i or later. Do not use Data Recovery with vCenter Servers running in linked mode.

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. While CIFS is supported it may not perform as well as VMDKs or RDMs and it is not recommended to use CIFS shares that are:

- On a server that has another role such as CIFS shares on a vCenter Server.
- Connected to a virtual machine.
- Shared to multiple services or servers.

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

Deduplication Store Sizing

The amount of storage required varies, depending on how much deduplication can save disk space as a result of running similar virtual machines. Data Recovery can attach to up to two deduplication stores for each backup appliance, and each deduplications store can be up to 1 TB in size. Data recovery is capable of protecting up to 100 virtual machines, though this number may be limited by disk space, depending on the size and complexity of virtual machines. Even with space savings, Data Recovery requires an absolute minimum of 10 GB of free space. This space is used for indexing and restore point processing, so even if the virtual machines to be backed up are very small, they may fail to complete if less that 10 GB of disk space is available. While a minimum of 10 GB is acceptable, having at least 50 GB is highly recommended for typical usage. 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.

For initial setup, provide storage space equal to the amount of used disk space on all virtual machines being protected. For example, if you are protecting 10 virtual machines, each with one 20 GB virtual disk each, and those virtual disks are on average 50% full, then you should provide at least 100 GB of storage available for the deduplication store. Over time, the amount of space the deduplication store consumes typically reaches an equilibrium as data being updated is roughly equal to aging restore points being removed by the retention policy.

Deduplication Store Formats

Deduplication stores can be stored on thin-provisioned or thick-provisioned virtual disks. Using thin-provisioning may result in decreased performance because space as allocated as it is required. Therefore, it may be best to use larger thick-provisioned disks sized to avoid the potential performance impact from growing a thin-provisioned disk. If the space available on a thick provisioned disk becomes unavailable, you can extend the disk using the vSphere Client.

Deduplication stores can be stored in all HCL supported storage and CIFS based network shares, and they are compatible with storage that is capable of deduplication. While any supported format may be used, virtual disks (VMDKs) or RDMs are recommended for deduplication stores because they provide the most well-understood and consistent performance. CIFS shares are also supported, but the performance of such shares varies across providers, and as such, is not an ideal solution. Furthermore, in many cases, virtual disks and RDMs perform better than network-based deduplication stores. Deduplication stores can be stored in RDM with either virtual or physical compatibility. If you plan to save the deduplication store to tape by taking snapshots, use an RDM with virtual compatibility. Snapshots cannot be taken with RDM with physical compatibility.

While CIFS can be used, do not use CIFS shares that are:

- On a server that has another role. For example, do not use CIFS shares hosted on a vCenter Server.
- Connected to a virtual machine.
- Shared to multiple services or servers.

NOTE Striping results in a loss of space efficiency across deduplication stores. Protecting virtual machines in separate deduplication stores typically provides better results than using striping to combine disks to create one large deduplication store.

Special Data Recovery Compatibility Considerations

There are special considerations to be aware of when establishing Data Recovery in your environment. Data Recovery does not support:

- vCenter Server in Linked Mode.
- IPv6 addresses. IPv4 addresses are required for the Data Recovery appliance.
- NFS is only supported if the share is presented by an ESX Server and the VMDK is assigned to the Data Recovery appliance.
- Hot adding disks with versions of vSphere that are not licensed for hot add.
- Restoring linked clones. Data Recover can backup linked clones, they are restored as unlinked clones.
- Backing up virtual machines that are protected by VMware Fault Tolerance.
- Backing up virtual machines that use VMware Workstation disk format.
- Backing up virtual machines with 3rd party multi-pathing enabled.
- Raw device mapped (RDM) disks in physical compatibility mode.

Data Recovery has been tested for use with:

- One backup appliance for each vCenter instance.
- Each backup appliance protecting up to 100 virtual machines.
- VMDK or CIFS based deduplication stores of up to 1TB.
- Up to two deduplication stores per backup appliance.

Install the Client Plug-in

Install the client plug-in on a computer that will be used to manage Data Recovery. You must install the client 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.

The client plug-in is only approved for managing backup appliances of the same version. Ensure you have the correct version of the plug-in for the appliance you are managing.

Procedure

1 Insert the Data Recovery installation CD.

  The VMware Data Recovery Installer window appears.

2 Click **Data Recovery Client Plug-In**.

3 Follow the prompts of the installation wizard.

4 Start the vSphere Client, and log in to a vCenter Server.

5 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 14.

Install the Backup Appliance

You must install the backup appliance on ESX 4.0 or later or ESXi 4.0 or later 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 vCenter Server and an ESX 4.0 or ESXi 4.0 host 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_OVF10.ovf` and select it.

The ovf file can be found on the Data Recovery CD in the `<Drive Letter>:\VMwareDataRecovery-ovf\` directory.

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 IP address allocation screen.

There are no configurable options on this screen. You can change IP address settings through the backup appliance console after installation. If such changes are required, use the vSphere Client to open the backup appliance console window, where you can modify IP address settings.

9 Select a timezone setting.

10 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 15. Otherwise you may now want to learn about Chapter 3, "Using VMware Data Recovery," on page 17.

If problems occur during deployment of the backup appliance, see http://www.vmware.com/support/developer/studio/studio20/va_user.pdf for more information about deploying virtual appliances.

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 Choose a type of storage.

- Select **Create a new virtual disk** and click **Next**.
- Select **Use an existing virtual disk** to add an existing disk such as when upgrading from older appliance and click **Next**.
- Select **Raw Device Mappings** to add the disk as an RDM and click **Next**.

7 If creating a new virtual disk, specify the disk size and other options and click **Next**.

8 If creating a new virtual disk, specify the advanced options and click **Next**.

9 Click **Finish**.

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, “Using VMware Data Recovery,” on page 17.

Using VMware Data Recovery 3

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

Common tasks involved with establishing and using 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

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 14. To help ensure timezone information is correct, when first powering on the backup appliance, use vCenter Server. After the first time the backup appliance is powered on, timezone information is set. After this information is set, the backup appliance can be powered on from the host without consequences to the timezone.

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.

The backup appliance is ready for use.

NOTE In vSphere Client under Inventory > Hosts and Clusters, the status for VMware Tools of the Data Recovery appliance status will indicate that it is Unmanaged. It is not necessary to update the VMware Tools on the Data Recovery appliance.

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 "Using Backup Jobs," on page 19 to create a backup job.

Using Backup Jobs

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

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 21.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 21.

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 restarts 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 18.

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 - Enter the IP address server name for the server hosting the network share.
- User name - The user name for an account with the required write privileges for the network share.
- Password - The password for the user account.

For information on adding a hard disk to the backup appliance, see "Add a Hard Disk to the Backup Appliance," on page 15.

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.

Bring to Compliance

You can make Data Recovery open the backup window for selected backup jobs until all applicable virtual machines are backed up. You may want to use this feature to create an initial set of backups after Data Recovery is first installed or to force all virtual machines backups to be made current. Virtual machines that have been backed up in the last 24 hours, regardless of how much they have changed since their last backup, are not backed up by Bring to Compliance.

Prerequisites

Before using the Bring to Compliance option, you must have installed and configured Data Recovery and you should have at least one backup job.

Procedure

1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.

2 Click the **Backup** tab, right-click a backup job, and click **Bring to Compliance**.

The backup window is held open so that backups can be performed on each virtual machine that has not been backed up in the last 24 hours. After these virtual machines are backed up, the backup window is returned to its previously defined configuration.

Override Backup Jobs

Backup job settings can be overridden so restore points are either kept by locking them or removed by marking them for deletion.

Prerequisites

Before you can lock restore points or mark them for removal, you must have installed and configured Data Recovery and you must have at least one restore point.

Procedure

1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.

2 Click the **Restore** tab, and select one or more restore points.

   a To mark restore points for deletion, click **Mark for Delete**.

   b To preserve restore points indefinitely, click **Lock**.

Restore points marked for deletion are deleted during the next integrity check or reclaim operation.

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.

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. The virtual machine created in the restore rehearsal has all NICs disconnected. This avoids the case where the trial restoration produces a virtual machine that starts completing tasks intended for an existing unrestored virtual machine.

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.

Understanding File Level Restore

Users may want to restore a version of a single file that was backed up using Data Recovery. Perhaps the file has been deleted or information from a previous version is required. In such a case, users can restore an entire previous version of the virtual machine that contained the file, but this may be cumbersome. Rolling back to previous versions may overwrite the existing virtual machine and even if the restored virtual machine is restored to an alternate location, the process may not be as fast as desired.

File Level Restore (FLR) addresses these issues by providing a way to access individual files within restore points for Windows virtual machines. This access makes it possible to read copies of files or restore them from within restore points to any other available location. For example, FLR makes it possible to create two copies of a file so the versions could be compared, or FLR could overwrite an existing file with an older version contained in the restore point, effectively reverting to a previous version.

Note that using FLR to access files in restore points only provides a way to read their contents. You cannot use FLR to modify the contents of a restore point. While FLR does not modify the contents of any restore points, some applications may make it appear that changes are occurring. For example, dragging and dropping a file from a restore point to another location may result in the file being removed from the list. This change does not reflect what has occurred. To confirm that the restore point contents have not been modified, refresh the view and note that all files are unchanged.

When FLR is invoked, all files required to complete file level restores are extracted from the executable. Similarly, when the FLR session ends, not only are all connections with restore points closed, but all files that were extracted are removed from the system.

When FLR starts, it establishes a connection with the Data Recovery backup appliance. FLR works in conjunction with VMware's Virtual Disk Development Kit (VDDK) to access information about the contents of restore points. All restore points are displayed, but FLR can only mount restore points for compatible Windows virtual machines.

When a restore point is mounted, a junction point is created on the virtual machine's local disk. The junction point is a directory that has the same name as the restore point. It contains a directory for each mounted disk associated with that restore point. Users can browse the contents of the vmdk disk files for the restore point for the virtual machine. Any files on the disk files for the selected restore point can then be copied to a location of the user's choosing.

After file level restore operations have been completed, you can choose to unmount individual restore points by selecting a restore point and click **Unmount**, or you can choose to unmount all restore points by clicking **Unmount All**.

After exiting FLR, all files that were extracted to enable FLR functionality are removed. Note that if FLR exits unexpectedly, previously extracted files are not removed. The number and size of files is not significant, so if an unexpected exit occurs, it is not imperative to clean up the files that remain. Upon the next use of the executable, any files that were left behind are used, and when subsequent sessions are ended, these residual files are removed from the system.

Install FLR

Install FLR on a Windows XP or later virtual machine by copying the FLR executable to a virtual machine.

Procedure

1 Insert the Data Recovery installation CD.

The VMware Data Recovery Installer window appears.

2 Click **Explore Media**.

3 Copy the FLR client executable from the installation CD at `<Drive Letter>:/WinFLR/VMwareRestoreClient.exe` to the virtual machine that will use the FLR client.

The FLR client is now ready for use on the Windows virtual machine.

Restore Files Using FLR Standard Mode

Use the File Level Restore (FLR) client to access individual files from restore points, rather than restoring entire virtual machines. This client is not required for the proper functioning of Data Recovery, but it does provide access to additional features.

Prerequisites

The FLR client can be used by users with Administrator privileges in virtual machines running Windows XP or later. The FLR client requires the .NET 2.0 framework. For FLR to be relevant, it is valuable to have a backup appliance with restore points. FLR can be installed to an environment that does not have a backup appliance or restore points, but without those things, the client will not be useful. In standard mode, files can only be restored for the virtual machine you are logged in to. FLR does not work with restore points for virtual machines that use GUID partition tables (GPT).

Procedure

1 Start the Windows virtual machine in which you will use FLR.

2 Double-click the FLR executable.

The VMware Data Recovery Restore Client window opens.

3 In the **IP address / Name** drop-down, select a Data Recovery appliance or enter the name or IP address of the appliance to which to connect and click **Login**.

FLR displays a list of all available restore points for the current virtual machine.

4 Select a restore point and click **Mount**.

The selected restore point is mounted as a directory on the local disk of the virtual machine being used. The contents of the restore point are now available and can be browsed from the virtual machine.

5 Browse or restore any desired files from the virtual machine.

6 When finished browsing or restoring files, click **Unmount All** and quit FLR.

Restore Files Using FLR Advanced Mode

Use FLR in advanced mode to access files from restore points from multiple virtual machines.

Prerequisites

The FLR client can be used by users with Administrator privileges in virtual machines running Windows XP or later. The FLR client requires the .NET 2.0 framework. For FLR to be relevant, it is valuable to have a backup appliance with restore points. FLR can be installed to an environment that does not have a backup appliance or restore points, but without those things, the client will not be useful. In advanced mode, files can be restored for any virtual machine that has been backed up. You must have access to an account with administrative privileges in vCenter to use advanced mode. FLR does not work with restore points for virtual machines that use GUID partition tables (GPT).

Procedure

1 Start the Windows virtual machine in which you will use FLR.

2 Double-click the FLR executable.

The VMware Data Recovery Restore Client window opens.

3 Select the **Advanced Mode** checkbox.

4 Provide FLR connection information.

a Under Data Recovery Appliance, in the **IP address / Name** drop-down, select a Data Recovery appliance or enter the name or IP address of the appliance to which to connect.

b Under vCenter Server, in the **IP address / Name** drop-down, select a Data Recovery appliance or enter the name or IP address of the appliance to which to connect.

c Under vCenter Server, in **User name** enter the name of a user with vCenter administrative privileges.

d Under vCenter Server, in **Password** enter the password for the previously specified administrative user.

e Click **Login**.

FLR displays a list of all available restore points for any backed up virtual machines on the Data Recovery appliance to which you are connected.

5 Select a restore point and click **Mount**.

The selected restore point is mounted as a directory on the local disk of the virtual machine being used. The contents of the restore point are now available and can be browsed from the virtual machine.

6 Browse or restore any desired files from the virtual machine.

7 When finished browsing or restoring files, click **Unmount All** and quit FLR.

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.
Data Recovery fails to complete backups with the error `disk full error –1115`, but the disk is not full.	Data Recovery requires disk space for indexing and processing restore points. As a result, Data Recovery typically needs enough free space to accommodate the size of the virtual machine backups plus an additional 10 GB. For example, to create a restore point for a single 10 GB virtual machine, a total of 20 GB should be available. To resolve this issue, add additional hard disks to the backup appliance.
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.
Data Recovery backup appliance is unmanaged.	This behavior is expected. The backup appliance is not managed by vSphere Server or other services such as Update Manager. It is not necessary and may not be possible to manage the backup appliance.

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. For more information on executing the log gathering script, see http://kb.vmware.com/kb/1012282.

You may also choose to review the verbose Data Recovery logs to determine if any helpful information is available there.

Understanding Damaged Restore Points

Restore points can become damaged due to storage medium failures and read/write errors. If such damage occurs, remove affected restore points.

Damaged restore points are identified during an integrity check. Any damaged restore points should be removed as they may block Data Recovery processes such as grooming. Review the Operations Log to find entries that refer to damaged restore points. If the log indicates that there are damaged restore points in your environment, remove them by either finding them in the inventory or finding all damaged restore points. After damaged restore points have been marked for deletion, run another integrity check to complete the process.

Remove Damaged Restore Points

Corrupt restore points, which are identified during integrity checks, should be removed. Restore points may be identified as damaged during transient connection failures. If transient connection failures are possible, check if damaged restore point issues are resolved after connections are restored.

Prerequisites

Before you can remove damaged restore points, you must have restore points in a functioning Data Recovery deployment.

Procedure

1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.

2 Click the **Reports** tab and double-click the integrity check that failed.

The Operations Log for the event opens in a separate window. Note which restore points triggered the failure.

3 Close the Operations Log and click the **Restore** tab.

4 From the Filter dropdown list, select **Damaged Restore Points**.

Available restore points are filtered to display only the virtual machines with damaged restore points. It may be necessary to expand a virtual machine's node to display the damaged restore point.

5 Select damaged restore points for removal and click **Mark for Delete**.

6 Initiate an integrity check.

Completing an integrity check causes all restore points marked for deletion to be removed.

7 Review the results of the integrity check to ensure no damaged restore points remain.

Understanding the datarecovery.ini File

The settings in the datarecovery.ini file can be modified to affect how the backup appliance completes tasks. Modifying the datarecovery.ini file is an advanced procedure that is typically use to change Data Recovery behavior in an attempt to troubleshoot problems.

Modify Backup Appliance Behavior Using the datarecovery.ini File

Making changes to the datarecovery.ini file affects the way the Data Recovery backup appliance behaves.

To complete this task, you will need access to an account with administrative permissions on the backup appliance.

Prerequisites

Before completing this procedure, the backup appliance must be powered on.

Procedure

1 Right-click the backup appliance virtual machine and choose **Open Console**.

2 Provide the username and credentials for this system.

It is recommended that the default username and password be changed as soon as the backup appliance is installed. If this was not changed, the default credentials are username: root, password: vmw@re.

3 Stop the datarecovery service using the command `service datarecovery stop`.

4 Using an editor of your choice, modify the datarecovery.ini file. If the datarecovery.ini file does not exist, create a file called `datarecovery.ini` in `/var/vmware/datarecovery`.

If you are creating a new datarecovery.ini file, the first line in the file must be `[Options]`. The datarecovery.ini file is case sensitive.

5 Save any changes and close the datarecovery.ini file.

6 Restart the datarecovery service using the command `service datarecovery start`.

datarecovery.ini Reference

Modify the settings in the ini file to affect the way that Data Recovery operates.

The contents of the datarecovery.ini file are case-sensitive.

Table 3-2. datarecovery.ini Settings

Option	Description	Example	Range	Default
MaxLogFiles	Sets the maximum number of log files that Data Recovery keeps. When the maximum is reached, the next created log file replaces the oldest existing log file.	MaxLogFiles=20		9
DisableHotaddCopy	Disables SCSI Hot-Add when set to 1.	DisableHotaddCopy=1	0-1.	0
DisableNetworkCopy	Disables network copy when set to 1.	DisableNetworkCopy=1	0-1.	0

Table 3-2. datarecovery.ini Settings (Continued)

Option	Description	Example	Range	Default
SetVCBLogging	The internal logging level for the VMware Consolidated Backup API.	SetVCBLogging=7	0-7. 7 is most verbose.	3
SetRAPILogging	The internal logging level for the Data Recovery API.	SetRAPILogging=7	0-7. 7 is most verbose.	3
SetEngineLogging	The internal logging level for the Data Recovery backup appliance.	SetEngineLogging=7	0-7. 7 is most verbose.	3
SetDevicesLogging	The internal logging level for the deduplication process.	SetDevicesLogging=7	0-7. 7 is most verbose.	3
SetAppLogging	The internal logging level for basic application logic.	SetAppLogging=7	0-7. 7 is most verbose.	3
SetVolumesLogging	The internal logging level for interactions between virtual machines and volumes.	SetVolumesLogging=7	0-7. 7 is most verbose.	3
SetBackupSetsLogging	The internal logging level for catalog operations.	SetBackupSetsLogging=7	0-7. 7 is most verbose.	3
IntegrityCheckInterval	The number of days between integrity checks.	IntegrityCheckInterval=7	0-7.	1
BackupRetryInterval	The number of minutes the backup appliance wait before retrying a failed backup.	BackupRetryInterval=20		30
RetentionPolicyInterval	The number of days before reclaiming space in the deduplication store.	RetentionPolicyInterval=4	1-7.	1
DedupeCheckOnRecatalog	Completes an integrity check after a recatalog when set to 1.	DedupeCheckOnRecatalog=1	0-1.	0

Table 3-2. datarecovery.ini Settings (Continued)

Option	Description	Example	Range	Default
EnableFileRestore	Disables File Level Restore when set to 0. This option only has an effect on Data Recovery version 1.1 or later. This option is ignored when FLR is used in Administrator Mode.	EnableFileRestore =1	0-1.	1
MaxBackupRestore Tasks	The maximum number of simultaneous backup and restores.	MaxBackupRestor eTasks=4	1-8.	8

Using Data Recovery Logs

Data Recovery provides logging that can vary in degree of detail and conditions under which it can be used.

Three notable types of logging include:

- Basic Logs - These logs provide basic information.
- Verbose Data Recovery Logs - These logs provide more extensive information.
- Client Connection Logs - These logs can be viewed even if you cannot connect to a backup appliance.

It is possible to view the logs for a single backup appliance. To review all logging information in an environment with multiple appliances, it is necessary to connect to each appliance and review that appliance's logs.

View the Data Recovery Logs

View the Data Recovery logs to gather information about the way the system is performing.

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 **Configuration** tab and click the **Log** link.

View the Verbose Data Recovery Logs

View the verbose data recover logs to find additional information about any issues you may be encountering.

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 **Configuration** tab and holding down the Shift key, click the **Log** link.

The Verbose log interface is displayed.

4 Click **Client Log**, **Appliance Operations Log**, or **Appliance Assert Log**, depending on the information you require.

5 To modify the logging level, hold down the Shift key and click **Refresh Log**.

The Logging Level control is displayed.

6 Click the up or down arrows on **Logging Level** to override the default settings.

View the Client Connection Logs

You can view the contents of the client connection logs, even if unable to connect to a backup appliance. The information in these logs may help solve connectivity issues.

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**.

If the connection succeeds, you can view normal logs, as described in "View the Data Recovery Logs," on page 31. If the connection fails, continue with this procedure.

3 Enter the keystroke series Ctrl-Alt-g-g.

The client connection logs are displayed.

Index