

VMware Data Recovery Administration Guide

Data Recovery 2.0

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-000665-00

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

Copyright © 2008–2011 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 is a registered trademark or trademark 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 Quiescing	8
Deduplication Store Benefits	10
2 Installing VMware Data Recovery	13
VMware Data Recovery System Requirements	13
Install the Client Plug-in	17
Install the Backup Appliance	17
Add a Hard Disk to the Backup Appliance	18
Extend a Disk	19
3 Using VMware Data Recovery	21
Understanding the Data Recovery User Interface	21
Power On the Backup Appliance	23
Configure the Backup Appliance	24
Connect the Backup Appliance to vCenter Server	25
Use the Getting Started Wizard	25
Using Backup Jobs	26
Establish a Maintenance Schedule	29
Configure Email Reporting	30
Restoring Virtual Machines	31
Understanding File Level Restore	33
Troubleshooting VMware Data Recovery	38
Index	47

About This Book

The *VMware Data Recovery Administration 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/ESXi™. 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 Quiescing,”](#) on page 8
- [“Deduplication Store Benefits,”](#) on page 10

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.

For virtual machines created in vSphere 4.0 or later, 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/ESXi 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/ESXi 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 apply to virtual machines created with hardware version 7 or later, but they do not apply to virtual machines created with VMware products prior to vSphere 4.0. For example, change block tracking is not used with virtual machines created with Virtual Infrastructure 3.5 or earlier. As a result, virtual machines created with earlier Hardware versions take longer to back up.

Data Recovery increases the speed and reduces the size of backups:

- 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 using the same job.
- Swap files are not backed up. This means that in Windows virtual machines, `pagefile.sys` file is not backed up, and in Linux, the swap partition is not backed up. This data is omitted since it is not relevant to restoring the system, thereby allowing backups to complete more quickly and to consume less disk space.

Data Recovery uses the vSphere licensing infrastructure to ensure that all virtual machines that are protected by Data Recovery have appropriate licensing.

Each instance of vCenter Server can support up to ten Data Recovery backup appliances and each 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 Quiescing

VMware Data Recovery uses Microsoft Windows Volume Shadow Copy Service (VSS) quiescing, 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 is instantiated by the VMware Tools service when a backup process is initialized. The VSP is registered as a Windows service and notifies the ESX/ESXi host when the applications are quiesced so it can take a snapshot of the virtual machine.

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

Table 1-1. Driver Type and Quiescing Mechanisms Used According to Guest Operating Systems

Guest Operating System	Driver Type Used	Quiescing Type Used
Windows XP 32-bit Windows 2000 32-bit	Sync Driver	File-system consistent quiescing
Windows Vista 32-bit/64-bit Windows 7 32-bit/64-bit	VMware VSS component	File-system consistent quiescing
Windows 2003 32-bit/64-bit	VMware VSS component	Application-consistent quiescing
On pre-ESX 4.1 Hosts: Windows 2008 32-bit/64-bit Windows 2008 R2	VMware VSS component	File-system consistent quiescing

Table 1-1. Driver Type and Quiescing Mechanisms Used According to Guest Operating Systems (Continued)

Guest Operating System	Driver Type Used	Quiescing Type Used
On ESX 4.1 and later Hosts: Windows 2008 32-bit/64-bit Windows 2008 R2	VMware VSS component	Application-consistent quiescing. For application-consistent quiescing to be available, three conditions must be met: <ul style="list-style-type: none"> ■ The UUID attribute must be enabled. This is enabled by default on virtual machines created on ESX 4.1 and later hosts. For virtual machines created on other hosts, complete the procedure “Enable Windows 2008 Virtual Machine Application Consistent Quiescing,” on page 40. ■ The virtual machine must use only SCSI disks. For example, application-consistent quiescing is not supported for virtual machines with IDE disks. There must as many free SCSI slots in the virtual machine as the number of disks. For example, if there are 8 SCSI disks on SCSI adapter 1, there are not enough SCSI slots free to perform application quiescing. ■ The virtual machine must not use dynamic disks.
Other guest operating systems	Not applicable	Crash-consistent quiescing

Application consistent quiescing of Windows 2008 virtual machines is only available when those virtual machines are created in vSphere 4.1 and later. Virtual machines created in vSphere 4.0 can be updated to enable application consistent quiescing, as described in “[Enable Windows 2008 Virtual Machine Application Consistent Quiescing](#),” on page 40.

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.

Detailed information about VSS can be found at <http://technet.microsoft.com/en-us/library/cc785914.aspx>.

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.

Deduplication Store Benefits

The deduplication store technology used by VMware Data Recovery 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 using 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 limited to using two deduplication stores. Data Recovery does not impose limits on the size of deduplication stores, but if the size of a store exceeds one terabyte, performance may be affected. While Data Recovery does not impose limits on deduplication store size, other factors limit deduplication shares. As a result, deduplication stores are limited to:

- 500 GB on CIFS network shares
- 1 TB on VMDKs and RDMs

NOTE NFS is only supported as a deduplication store format if the share is presented by an ESX/ESXi Server and the VMDK is assigned to the Data Recovery appliance.

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. Integrity checks are completed on some or all of the deduplication store under different conditions. Data Recovery is designed to complete an incremental or full integrity checks during the maintenance window. Incremental integrity checks verify the integrity of restore points that have been added to the deduplication store since the most recent full or incremental integrity check. Data Recovery is also designed to perform an integrity check of all restore points once a week.

The destination maintenance window should be used to avoid the case where integrity checks may consume computing resources or otherwise interfere with any backup operations in process. As a result, the destination maintenance window and backup window should be defined such that they do not overlap. The destination maintenance is stopped if it does not complete within the defined window. Even if the maintenance is stopped, the destination is not locked out from other operations such as backup and restore. The next time destination maintenance window opens, the operation continues where it was left off. For more information on configuring the maintenance window, see [“Establish a Maintenance Schedule,”](#) on page 29.

In addition, the integrity check can be started manually. When the integrity check is started manually, it always performs full integrity check of the entire destination, and does not use the maintenance window. 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 restores are not allowed.

Data Recovery stores information about the progress of an integrity check. As a result, if the backup appliance stops integrity check, the process can be restarted from where the check was stopped, thereby ensuring that work completed on an integrity check is not lost. The backup appliance stops integrity checks when the maintenance window passes. Tracking progress helps ensure integrity checks eventually complete. Integrity checks that are manually stopped by user intervention do not save progress information, so after such a stop, the integrity check begins again from the start.

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 according to the maintenance window. While the reclaim operation is in progress, backups to the deduplication store are not allowed, but restore operations from the deduplication store are allowed.

The reclaim operation starts or is deferred based on the same logic used for determining whether or not to complete an integrity check. Reclaim operations are generally run once every 24 hours when no backup windows are active.

During the reclaim operation, Data Recovery applies the retention policy for each source virtual machine in a backup job for the corresponding destination. If one virtual machine is included in multiple backup jobs with different retention policies, Data Recovery combines the retention policy, keeping sufficient backups to meet the criteria of all backup jobs. If a source virtual machine was defined in a backup job at some point, but the virtual machine is deleted or is no longer defined in a backup job, none of the restore points of that virtual machine are removed.

The retention policy keeps backups that are some combination of being weekly, monthly, quarterly, or yearly. Those periods are defined as follows:

Table 1-3. Criteria for Determining Different Types of Backups

Backup Type	Criteria
Weekly	The first backup after 10:00 PM on Friday.
Monthly	The first backup after 10:00 PM on the last day of the month.
Quarterly	The first backup after 10:00 PM on the last day of the month for March, June, September, and December.
Yearly	The first backup after 10:00 PM on December 31st.

NOTE When reclaim operations free space in files, those files are not compacted to reflect the new free space. As a result, the amount of free space on the deduplication store does not increase, even when reclaim operations are reclaiming space. The space which is free is reserved and used for future backups.

Installing VMware Data Recovery

VMware Data Recovery uses a plug-in to the vSphere Client and a backup appliance to store backups to 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/ESXi 4 host.
- The optional File Level Restore (FLR) client is installed in a virtual machine running a supported guest operating system. For more information on FLR, see [“Understanding File Level Restore,”](#) on page 33.

This chapter includes the following topics:

- [“VMware Data Recovery System Requirements,”](#) on page 13
- [“Install the Client Plug-in,”](#) on page 17
- [“Install the Backup Appliance,”](#) on page 17
- [“Add a Hard Disk to the Backup Appliance,”](#) on page 18
- [“Extend a Disk,”](#) on page 19

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/ESXi 4 or later. The ESX/ESXi host that runs the backup appliance must be managed by vCenter Server.
- When using Data Recovery with vCenter Servers running in linked mode, login to the vCenter Server with which the Data Recovery appliance is associated.

You can store backups on any virtual disk supported by ESX/ESXi. 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.

When adding hard disks to the backup appliance, consider how many disks most virtual machines to be backed up have. Each backup appliance can back up 100 virtual machines, but a maximum of 8 virtual machines can be backed up simultaneously. Each disk on each virtual machine may be hot-added for the backup to occur. In the default configuration, the backup appliance has a SCSI adapter #0 and a SCSI disk #0 attached to the SCSI adapter. Since the first SCSI adapter has a system disk at SCSI 0:0, only 14 SCSI disks can be hot-added. As the backup of a virtual machine completes, that virtual machines disks are removed and subsequent backups can begin. In the default configuration, if the total number of disks for the virtual machines being backed up reaches 15, the disks are backed up over the network instead of through hot-add. If you are working with virtual machines with a greater number of disks, consider adding additional disks to the appliance. For example, if each virtual machine in your environment has 3 disks, some of the virtual machine disks are backed up over the network, and performance may be negatively affected. By adding a dummy disk of 1 MB to another SCSI bus adapter, the total available SCSI bus locations for hot-adding increases to 30, so all 8 virtual machines in the example given here can be backed up simultaneously using hot-add. Additional disks should be added in the sequence SCSI 1:0, SCSI 2:0, SCSI 3:0, and so on for as many or as few instances as are required. A virtual machine, such as the backup appliance may have up to 4 SCSI adapters, enabling a maximum of 60 available SCSI bus locations for hot-adding disks, which is sufficient for most environments.

See the most recent vSphere documentation for information about setting up a vSphere 4.0 or later 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. 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 than 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, 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 is 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.

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. If multiple appliances use a single CIFS share, miscalculations in space requirements may result, which may cause the appliance to run out of disk space.

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.

There are special considerations when using thin provisioned virtual disks as the data recovery destination disk. vSphere automatically freezes any virtual machine whose thin provisioned disk usage exceeds its hosting VMFS datastore's capacity. Therefore, VMware recommends using one of two strategies to avoid running out of space for the Data Recovery destination disk.

- Use alarms to identify when space is limited on a thin-provisioned disk and add more space as required.
- Use smaller thick provisioned virtual disks and extend the disk as required.

Networking Requirements

Different components of Data Recovery communicate among each other over TCP. As a result, ensure the appropriate ports are open in your environment for normal operation.

- The backup appliance connects to vCenter Server web services. By default, this connection is established using ports 80 and 443.
- The Data Recovery client plug-in and File Level Restore (FLR) client connect to the backup appliance using port 22024.
- The backup appliance connects to VMware ESX or VMware ESXi using port 902.

ESX/ESXi servers that were added to vCenter using a DNS name must have a name that is resolvable. In some cases, using DNS names creates problems. If problems arise with resolving DNS names, consider adding ESX/ESXi servers using IP addresses instead.

Security Credentials Requirements

Data Recovery completes operations using permissions that are granted through a role that is assigned to a user. The specific permissions that must be assigned to roles vary based on the task to be completed.

If your backup infrastructure uses Network Block Device (NBD) technology, the following minimum permissions must be assigned to the role that the appliance runs as:

- VirtualMachine->Configuration->Disk change tracking
- VirtualMachine->Provisioning->Allow read-only disk access
- VirtualMachine->Provisioning->Allow virtual machine download
- VirtualMachine->State->Create snapshot
- VirtualMachine->State->Remove snapshot
- Global>DisableMethods
- Global>EnableMethods
- Global->License

If your backup infrastructure uses SCSI hot-adding, the role that the appliance runs as must have all permissions required by NBD, as well as the following additional minimum permissions:

- Datastore->Allocate space
- VirtualMachine->Configuration->Add existing disk

- VirtualMachine->Configuration->Add new disk
- VirtualMachine->Configuration->Add or remove device
- VirtualMachine->Configuration->Change resource
- VirtualMachine->Configuration->Remove disk
- VirtualMachine->Configuration->Settings

Special Data Recovery Compatibility Considerations

There are special considerations to be aware of when establishing Data Recovery in your environment. Data Recovery is supported for use with:

- Ten Data Recovery backup appliances for each vCenter Server instance.
- Each backup appliance protecting up to 100 virtual machines.
- VMDK or RDM based deduplication stores of up to 1TB or CIFS based deduplication stores of up to 500 GB.
- CIFS shares with passwords limited to 64 characters or less. CIFS share passwords must conform to the Latin 1(ISO 8859-1) standard. Double-byte characters are not supported.
- If a third-party solution is being used to backup the deduplication store, those backups must not run while the Data Recovery service is running. Do not back up the deduplication store without first powering off the Data Recovery Backup Appliance or stopping the data recovery service using the command `service datarecovery stop`.
- Up to two deduplication stores per backup appliance.
- vCenter Servers running in linked mode. For this configuration to perform as expected, log in to the vCenter Server with which the Data Recovery appliance is associated.

Data Recovery does not support:

- IPv6 addresses. IPv4 addresses are required for the Data Recovery appliance.
- Hot adding disks with versions of vSphere that are not licensed for hot plug.
- Restoring VMware View linked clones. Data Recovery can back up VMware View linked clones, but 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 where shared SCSI buses are in use.
- Raw device mapped (RDM) disks in physical compatibility mode in virtual machines to be backed up.
- Using older versions of the vSphere Client plug-in or older versions of FLR with the current version of Data Recovery.
- Multiple backup appliances on a single host.
- Using Data Recovery to backup Data Recovery backup appliances. While this is not supported, this should not be an issue. The backup appliance is a stateless device, so there is not the same need to back it up as exists for other types of virtual machines.
- Backup of virtual machine disks that are marked as Independent.
- Backup of Storage Virtual Appliances (SVAs).

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 17.

Install the Backup Appliance

Install the backup appliance on ESX/ESXi 4.0 Update 2 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 you should have an ESX/ESXi 4.0 Update 2 host running in your environment. The backup appliance connects to ESX/ESXi using port 902. If there is a firewall between the backup appliance and ESX/ESXi, port 902 must be open. The backup appliance, client plug-in, and FLR should all be the same version. Do not install multiple backup appliances on a single host.

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 Select a location for the backup appliance in the vSphere inventory.
You can optionally rename the backup appliance.
- 5 Select the host or cluster to which the backup appliance is to be deployed.

- 6 Select a datastore to store the virtual machine files.
When choosing a datastore on which to store the files for the backup appliance, choose a datastore with the largest VMFS block size. This is necessary to ensure that the backup appliance can back up virtual machines from all datastores.
- 7 Select a disk format to use for the virtual disk.
- 8 In Properties, select a timezone for the appliance.
- 9 Review the deployment settings and click **Finish**.
- 10 If you choose to enable the second NIC provided with the backup appliance:
 - a Enter the command `ifup eth1` in the command line to enable eth1.
 - b Change `ONBOOT=no` to `ONBOOT=yes` in `/etc/sysconfig/network-scripts/ifcfg-eth1` to automatically enable eth1 when the appliance is started.

The backup appliance is now deployed into your environment.

What to do next

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.

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 18. Otherwise you may now want to read [Chapter 3, “Using VMware Data Recovery,”](#) on page 21.

The backup appliance is recognized by an annotation on the virtual machine that says VMware Data Recovery Module. Do not change this annotation or add this annotation to any other virtual machines. Manually adding or removing this annotation will produce undesirable results.

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. For more information on disk formats, including using thin provisioned disks see [“Deduplication Store Formats,”](#) on page 14. For more information on the value of adding SCSI disks, see [“VMware Data Recovery System Requirements,”](#) on page 13.

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 an 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**.
If creating a SCSI virtual disk, it is recommended that you set the SCSI value to SCSI 1:0.
- 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 may not be immediately recognized. Either wait until the hard disk appears or reboot backup appliance.

What to do next

You may now want to learn about [Chapter 3, “Using VMware Data Recovery,”](#) on page 21.

Extend a Disk

To make more space available, disks can be extended.

Prerequisites

Extending a disk requires that there be a disk available with free space to accommodate the extension.

Procedure

- 1 Start the vSphere Client and log in to the vCenter Server that manages the backup appliance.
- 2 Check to ensure no operations are currently being completed on the disk.
- 3 Select **Inventory > Hosts and Clusters**.
- 4 Select the appropriate hard drive and increase the size.

The disk is not extended, but it is necessary to wait for a few minutes for the operating system to recognize the updated disk configuration. If the disk is not recognized, consider rebooting the backup appliance.

Using VMware Data Recovery

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:

- Configuring Data Recovery.
- Establishing backup jobs, including required resources, which may include adding network shares or formatting volumes.

When using Data Recovery with vCenter Servers running in linked mode, you must login to the vCenter Server with which the Data Recovery appliance is associated.

This chapter includes the following topics:

- [“Understanding the Data Recovery User Interface,”](#) on page 21
- [“Power On the Backup Appliance,”](#) on page 23
- [“Configure the Backup Appliance,”](#) on page 24
- [“Connect the Backup Appliance to vCenter Server,”](#) on page 25
- [“Use the Getting Started Wizard,”](#) on page 25
- [“Using Backup Jobs,”](#) on page 26
- [“Establish a Maintenance Schedule,”](#) on page 29
- [“Configure Email Reporting,”](#) on page 30
- [“Restoring Virtual Machines,”](#) on page 31
- [“Understanding File Level Restore,”](#) on page 33
- [“Troubleshooting VMware Data Recovery,”](#) on page 38

Understanding the Data Recovery User Interface




The vSphere Client plug-in for Data Recovery provides a number of new user interface elements that can be used for configuring Data Recovery behavior.

The Data Recovery user interface is divided into several tabs. Tabs with new interface options include: the Getting Started tab, the Backup tab, and the Restore tab.

Getting Started tab

The Getting Started tab provides introductory information about Data Recovery and provides a way to start common configuration tasks.




Table 3-1. Getting Started tab

Icon	Name	Description
	Add a Job	Launches the Backup Job wizard. For more information, see “Use the Backup Job Wizard,” on page 27.
	Restore a Virtual Machine	Launches the Restore a Virtual Machine wizard. For more information, see “Restore Virtual Machines from Backup,” on page 32.
	View Reports	Switches the current view to the Reports tab, which provides a way to review the status of existing jobs.

Backup tab

The Backup tab displays information about existing backup jobs and their status and provides a way to create, edit, and delete backup jobs.





Table 3-2. Backup tab

Icon	Name	Description
	Add a Job	Launches the Backup Job wizard. For more information, see “Use the Backup Job Wizard,” on page 27.
	Edit a Job	Launches the Backup Job wizard for editing an existing job.
	Delete a Job	Deletes the selected backup job.

Restore tab

Existing restore points can be restored, locked, or marked for delete in the Restore tab. The process of locking and marking for deletion are mutually exclusive, so you can only select one of those two options. For more information on locking restore points or marking restore points for deletion, see [“Mark Restore Points for Removal or Locking,”](#) on page 29. The Restore tab may be unavailable if there are no existing restore points.

Table 3-3. Restore tab

Icon	Name	Description
	Restore a Virtual Machine	<p>Launches the Restore Virtual Machines from Backup, which provides a way to configure how virtual machines are restored to the state saved in the selected restore points. For more information, see “Restoring Virtual Machines,” on page 31.</p> <p>By default, Data Recovery manages the storage and eventual deletion of older restore points according to the Retention Policy specified in the backup job. The icon for restore points being managed by Data Recovery appear as follows: </p>
	Lock a Restore Point	Any selected restore points are toggled between being locked or unlocked. Locked restore points are preserved indefinitely rather than being eliminated over time according to the Retention Policy.
	Delete a Restore Point	Any selected restore points are toggled between being marked for delete or not marked for delete. Restore points that are marked for deletion are removed by Data Recovery processes. Restore points marked for deletion are typically not deleted immediately.

Power On the Backup Appliance

The virtual machine backup appliance must be powered on to perform backups. The backup appliance is automatically powered on in some cases, but you may choose to power the backup appliance on manually, for example, to change the password.

Prerequisites

Before powering on the backup appliance, you must [“Install the Client Plug-in,”](#) on page 17 and [“Install the Backup Appliance,”](#) on page 17. Using mismatched versions of the plug-in is not supported and may result in errors that indicate that the backup appliance is not powered on.

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.

What to do next

If you need to shut down or restart the backup appliance, do not do so while backups are in progress. Before shutting down the appliance, stop all backups using the Data Recovery client, wait for the backups to stop, and then shut down the appliance.

Configure the Backup Appliance

You can configure networking settings or reboot the backup appliance, as required, using the web interface. If the backup appliance was deployed through vCenter Server, the backup appliance timezone is automatically configured. If the backup appliance was installed through and ESX/ESXi server, it may be necessary to configure timezone information.

Prerequisites

Before you can configure the backup appliance, it must be powered on and the current version of the client plug-in should be installed.

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 not being managed by vSphere. Do not update the VMware Tools on the Data Recovery appliance. The unmanaged status means that the appliance is not being managed by vSphere, but it is being managed by Data Recovery.

Connect the Backup Appliance to vCenter Server

The VMware Data Recovery backup appliance must be connected to the vCenter Server to perform automated tasks such as automated backups and reclaim operations. Before connecting the backup appliance, it must be powered on.

Prerequisites

Typically, when a backup appliance is powered on, it is automatically connected to a vCenter Server, but you may need to complete this process manually. To connect the backup appliance, you can use either the virtual machine name or the IP address. Using a name requires a name resolution service and unique backup appliance name. If your environment does not include a name resolution service or has multiple backup appliances with the same name, the connection might fail. In such a case, enter the IP address and try again.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 Select the backup appliance from the inventory list in the left pane. Backup appliance names are displayed in bold to help identify possible choices. Alternately, you can enter the virtual machine name or IP address of the backup appliance. Click **Connect**.
 - ◆ If this is the first time a vSphere Client has connected to the backup appliance, the Getting Started wizard is launched automatically. Complete the wizard, as described in [“Use the Getting Started Wizard,”](#) on page 25.
- 3 In the **Configuration** tab, select **Backup Appliance**.
- 4 Click the **Set vCenter Server or ESX/ESXi Host** link.
- 5 Enter the vCenter user name and password and click **Apply**. The appliance stores the information required to connect to vCenter Server to perform backup and restore operations.

The backup appliance is now connected to the vCenter Server and backups can now be completed.

What to do next

Next you may choose to create backup jobs as described in [“Use the Getting Started Wizard,”](#) on page 25 or [“Using Backup Jobs,”](#) on page 26.

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,”](#) on page 25. The Getting Started Wizard starts automatically after the first time connecting to the backup appliance, in which case, begin with [Step 4](#).

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 If this is not the first time connecting to the backup appliance, start the Getting Started Wizard by clicking the **Configuration** tab and clicking **Getting Started Wizard**
- 3 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.

- 4 In the Backup Destinations page, select a backup destination from the list of choices.
- 5 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 formatted disk, 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. The CIFS share password is limited to 64 characters or less and must conform to the Latin 1(ISO 8859-1) standard. Double-byte characters are not supported.
- 6 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 26 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.

Data Recovery uses the backup window to create new backups and the retention policy to remove specific older ones. For more information on how the deduplication store processes of integrity checks and reclaim operations support this functionality, see [“Deduplication Store Benefits,”](#) on page 10.

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.

NOTE Using Data Recovery to back up the Data Recovery backup appliance is not supported.

Destination

You can store backups in VMDKs, on RDMs, or on network shares. 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 28. You must format VMDKs and RDMs to store backups. You can format destinations that are not yet formatted or partitioned. For more information, see [“Formatting a Volume,”](#) on page 28.

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 backup 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. For more information on how different backup periods are assessed, see [Table 1-3](#).

If the deduplication store is less than 80% full, the retention policy is run once each week. If the deduplication store is more than 80% full, the retention policy is run once each day. If the deduplication store is full, the retention policy is run immediately if the policy has not been run within the last 12 hours.

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 25.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery** and click **Connect**.
- 2 Click the **Backup** tab and click **New** to launch the Backup Job wizard.
- 3 In the Name page, accept the suggested name or enter an alternate name and click **Next**.
- 4 In the Virtual Machines page, select individual virtual machines or containers that contain virtual machines to be backed up and click **Next**.
- 5 In the Destinations page, select a storage destination and click **Next**.
- 6 In the Backup Window page, accept the default times or specify alternate backup windows and click **Next**.
- 7 In the Retention Policy page, accept the default retention policy or specify an alternate retention policy and click **Next**.
- 8 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. For example, a valid URL might be \\192.168.12.1\C\$ or \\MyNetworkShare\MySharedDirectory.

NOTE Adding network shares is only supported at the share level. Attempts to mount subfolders on a share will result in the root level being mounted.

- User name - The user name for an account with the required write privileges for the network share.
- Password - The password for the user account. Older versions of VMware Data Recovery may restrict password length and use of non-ASCII characters.

For information on adding a hard disk to the backup appliance, see [“Add a Hard Disk to the Backup Appliance,”](#) on page 18.

Formatting a Volume

VMware Data Recovery can store backups on VMDKs, RDMs, and network volumes. 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.

Backup Now

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 Backup Now.

Prerequisites

Before using the Backup Now 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** and click **Connect**.
- 2 Click the **Backup** tab, right-click a backup job, and click **Backup Now** and select either **All Sources** or **Out of Date Sources**.

Backups are run immediately and complete even if the backup window is closed. If you selected **All Sources**, all virtual machines are backed up. If you selected **Out of Date Sources**, all virtual machine that have not been backed up in the last 24 hours are backed up.

Suspend Backups

A backup job can be manually suspended, meaning that no new backup jobs are started.

Prerequisites

Before using the Suspend Backups 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** and click **Connect**.
- 2 Click the **Backup** tab, right-click a backup job, and click **Suspend Future Tasks**.

Backups will not be completed until this setting is reversed using the same process. The Suspend Future Tasks setting is not persistent, so backups resume if the backup appliance is restarted.

Mark Restore Points for Removal or Locking

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** and click **Connect**.
- 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. To force the immediate deletion of restore points, manually start an integrity check.

Establish a Maintenance Schedule

Data Recovery maintenance is completed through integrity checks and reclaim operations.

You can schedule the times at which these operations occur, allowing them to be completed during times when the system is likely to be idle. For example, you might configure the Maintenance Schedule window to not overlap with scheduled backups.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery** and click **Connect**.
- 2 Click the **Configuration** tab, click **Destinations**, and click **Maintenance Schedule...**

The Destination Maintenance Schedule window appears.

- 3 Select the days and times during which maintenance tasks are permitted, and then click **OK**.

For example, you might select all, select some subset, or click Default, which allows maintenance to be completed between 9 AM and 5 PM during weekdays.

After configuring a maintenance schedule, integrity checks and reclaim operations are completed during the specified hours. If the backup appliance is in the process of conducting an integrity check and the maintenance window closes, progress on the check is saved. When the maintenance window re-opens, the integrity check recommences from the point it was at when it was last stopped.

Configure Email Reporting

Information about Data Recovery operations can be emailed to specified users. This can be used to receive current information about the health of the Data Recovery system without having to connect to the Data Recovery appliance.

Data Recovery supports standard SMTP (port 25), authenticated SMTP (port 587), and non-standard ports in which you add the port number to the end of the mail server. If you do not specify a port, the defaults are used. Data Recovery email reporting does not support secure SMTP. Email reporting does not support sending messages protected with SSL encryption.

Prerequisites

To use this feature, a mail server must be available. You must have several pieces of information to properly configure email reporting. This includes:

- The outgoing mail server name. An example might be smtp.example.com.
- If the mail server requires authentication, the username and password for an account with the appropriate privileges.

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery** and click **Connect**.
- 2 Click the **Configuration** tab, click **Email**, and click **Properties...** The Email Properties window appears.
- 3 Specify settings for email reporting.
 - a Click **Enable email reports** to enable the feature.
 - b Provide the mail server name.
 - c Provide the user name and password for an account with the appropriate permissions, as required.
 - d Provide the address from which the mail will be sent.
This field cannot be empty.
 - e Provide the addresses to which the mail will be sent.
Up to ten addresses may be used. Multiple addresses must be separated by commas (,) or semi-colons (;). If an email address does not include an at symbol (@), the mail will be sent to the user at the domain name for the host server. For example, if the host mail server is smtp.example.com and the email address user is provided, an email will be sent to user@example.com.
 - f Specify at what time and on which days the email report will be sent.
- 4 You may elect to test the settings by clicking **Send test email**.
- 5 Click OK.

Email reporting is now configured. An email with information about the Data Recovery system state is sent at the configured times.

Restoring Virtual Machines

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

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 backup jobs, you can specify collections of virtual machines, such as all virtual machines in a datacenter. It is possible to move virtual machines and VMDK files to different locations. 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 restore 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.

Restore Virtual Machines from Backup

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

Prerequisites

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

Procedure

- 1 In the vSphere Client, select **Home > Solutions and Applications > VMware Data Recovery**.
- 2 Connect to the backup appliance.
- 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. If you did not connect to the backup appliance using an account that has been assigned the Administrator role, click **Restore Credentials** and provide a username and password for an account with that role assignment. 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.

Complete a Restore Rehearsal from Last Backup

Restore Rehearsal from Last Backup creates a new virtual machine from the most recent backup of the selected virtual machine. Complete a Restore Rehearsal from Last Backup to confirm that a virtual machine is being backed up as expected and that a successful restore operation can be completed.

Prerequisites

Before you can complete a Restore Rehearsal from Last Backup, 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 Connect to the backup appliance.
- 3 Right-click a virtual machine that has a backup and select **Restore Rehearsal from Last Backup**.

The Virtual Machine Restore wizard appears displaying the Sources page. The most recent backup of the VM that was selected in the inventory tree is selected by default.

- 4 Review the proposed settings on the Sources page. You may choose to modify the provided settings. Click **Next**.

The Destinations page appears.

- 5 Review the proposed settings on the Destinations page. You may choose to modify the provided settings. If you did not connect to the backup appliance using an account that has been assigned the Administrator role, click **Restore Credentials** and provide a username and password for an account with that role assignment. Click **Next**.

A new virtual machine with "Rehearsal" appended to its name is created by default in the same location as the source virtual machine. You may choose to rename the new virtual machine and change the location where it will be created on this page.

The Ready to Complete page appears

- 6 Click **Restore** to complete the restore rehearsal from last backup or click **Back** to modify settings.

A version of the virtual machine is restored to the inventory. The virtual machine created in the 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.

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 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.

Using FLR to access files in restore points only provides a way to read their contents. Do not attempt to 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. Similarly, it is possible to open the files contained in restore points, make changes, and save and close those files. This does not change the information stored in the restore point in the deduplication store. As a result, when users exit FLR, any changes that appeared to be made to files in a restore point are lost. To save such changes, either create and edit local copies outside of the restore point, or edit the contents of the restore point by starting the virtual machine and modifying the files in the virtual machine.

If the backup appliance is completing other tasks such as running backup or restore jobs, FLR may be delayed in establishing a connection. All restore points are displayed, but FLR can only mount restore points for compatible virtual machines. Some file systems may not be mountable by a specific virtual machine. FLR uses the operating system on which it is running to read the contents of restore points. As a result, if the operating system of the virtual machine in which FLR is running can not read the file system for the restore point, that restore point will be inaccessible. For example, Linux machines may be unable to read NTFS files, so attempting to use FLR in a Linux virtual machine to read the contents of a Windows virtual machine's restore point may fail.

When a restore point is mounted, a root mount point is created on the virtual machine's local disk. The root mount point is a directory that has the same name as the restore points date in long format. 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 restore points. To unmount individual restore points in Windows, select a restore point and clicking **Unmount**, or you can choose to unmount all restore points by clicking **Unmount All**. To unmount restore points using FLR on Linux, enter the command `unmount`.

After exiting FLR, all resources that were used to enable FLR functionality are removed. Note that if FLR exits while mount points are still busy, you may need to perform a manual clean up of these resources. For more information on manually cleaning up busy unmounts, see the release notes.

The FLR client can be used by users with Administrator privileges in Windows or sudo privileges in Linux virtual machines. In Windows virtual machines, the FLR client requires the .NET 2.0 framework or later. In Linux virtual machines, the FLR client requires the 32-bit version of FUSE 2.5 or later. Note that for Linux, the 32-bit version is required, regardless of whether the virtual machine being used is 32-bit or 64-bit. 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. Use matching versions of FLR and the backup appliance. Using an older version of FLR may fail. FLR does not work with restore points for virtual machines that use GUID partition tables (GPT). FLR can be installed and used on virtual machines running the following operating systems:

- 32 or 64-bit Linux virtual machines including:
 - Red Hat Enterprise Linux (RHEL) 5.4/CentOS 5.4
 - Red Hat 4.8/CentOS 4.8
 - Ubuntu 8.04
 - Ubuntu 8.10
 - Ubuntu 9.04
- Windows virtual machines including:
 - Windows XP
 - Windows Vista
 - Windows 7
 - Windows Server 2003
 - Windows Server 2008

NOTE FLR is not recommended for use in some cases.

- FLR is not supported on physical machines.
 - FLR is not recommended for use in environments using VMware vCloud Director (vCD). FLR and vCD may interact in unexpected ways, so disabling FLR is recommended by setting the `EnableFileRestore` to 0 in the `datarecovery.ini` file. For more information on modifying the `datarecovery.ini` file, see ["Understanding the datarecovery.ini File,"](#) on page 41.
-

Use FLR in Windows

Use FLR on a Windows virtual machine by copying the FLR executable to that 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 Windows virtual machine that will use the FLR client.

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

Restore Files Using FLR Standard Mode in Windows

Use the File Level Restore (FLR) client in a Windows virtual machine 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

Before restoring files, complete the steps described in [“Use FLR in Windows,”](#) on page 35. FLR connects to the backup appliance using port 22024. If there is a firewall between the FLR client and ESX/ESXi, port 22024 must be open before restore points can be accessed using FLR. To work with files on other virtual machines, use advanced mode, as described in [“Restore Files Using FLR Advanced Mode in Windows,”](#) on page 35.

Procedure

- 1 Start the 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 in Windows

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

Prerequisites

FLR connects to the backup appliance using port 22024. If there is a firewall between the FLR client and ESX/ESXi, port 22024 must be open before restore points can be accessed using FLR.

Procedure

- 1 Start the 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 the user used to connect to the backup appliance. This must be 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 With the mounted restore point selected, click **Browse** to open an instance of Windows Explorer at the location of the mounted files.
- 7 Browse or restore any desired files from the virtual machine.
- 8 When finished browsing or restoring files, click **Unmount All** and quit FLR.

Use FLR in Linux

Use FLR on a Linux virtual machine by copying the FLR executable to that virtual machine.

Prerequisites

In Linux virtual machines, the FLR client requires the 32-bit version of FUSE 2.5 or later be installed. This is a requirement for both 32-bit and 64-bit Linux virtual machines. On Linux, FLR requires fuser and LVM. FLR uses fuser during unmount attempts to determine if mounts are busy and uses LVM to access LVM volumes. To make these utilities available, they must be installed and added to PATH.

Procedure

- 1 Insert the Data Recovery installation CD.
- 2 Copy the FLR client archive `LinuxFLR/VMwareRestoreClient.tgz` on the installation CD to the virtual machine that will use the FLR client.
- 3 Extract the archive using `tar xzvf VMwareRestoreClient.tgz`.

- 4 Navigate to the `VMwareRestoreClient` directory and invoke FLR by executing `./VdrFileRestore`.

Ensure you use `VdrFileRestore` rather than `vdrFileRestore`. These are two separate executables.

`VdrFileRestore` is a wrapper script that includes `vdrFileRestore` and provides additional benefits such as setting up correct library dependencies and ensuring the proper FUSE installation is available.

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

Restore Files Using FLR Standard Mode in Linux

Use the File Level Restore (FLR) client in a Linux virtual machine 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. For a full list of command options available with `VdrFileRestore`, see the readme file included in the Linux FLR tgz file.

Prerequisites

Before restoring files, complete the steps described in [“Use FLR in Linux,”](#) on page 36. FLR connects to the backup appliance using port 22024. If there is a firewall between the FLR client and ESX/ESXi, port 22024 must be open before restore points can be accessed using FLR. To work with files on other virtual machines, use advanced mode, as described in [“Restore Files Using FLR Advanced Mode in Linux,”](#) on page 37.

Procedure

- 1 Start the virtual machine in which you will use FLR.
- 2 Execute `VdrFileRestore` by supplying an IP address or name of the Data Recovery appliance using the syntax (`-a | --appliance <ip | dns name>`). An example of this would be the command: `./VdrFileRestore -a 10.0.1.124`

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

- 3 Select a restore point.

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.

- 4 Browse or restore any desired files from the virtual machine.
- 5 When finished browsing or restoring files, enter the command `umount` and FLR exits.

Restore Files Using FLR Advanced Mode in Linux

Use FLR in a Linux virtual machine in advanced mode to access files from restore points from multiple virtual machines. For a full list of command options available with `VdrFileRestore`, see the readme file included in the Linux FLR tgz file.

Prerequisites

FLR connects to the backup appliance using port 22024. If there is a firewall between the FLR client and ESX/ESXi, port 22024 must be open before restore points can be accessed using FLR.

Procedure

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

- 2 Execute VdrFileRestore. At a minimum you must supply an IP address / Name of the Data Recovery appliance (`-a <ip | dns name>`), an IP address / Name of the vCenter Server (`-s <ip | dns name>`), a user name of a user with vCenter administrative privileges (`-u <user>`), and a password for the previously specified administrative user (`-p | --password <password>`). An example of this would be the command: `./VdrFileRestore -a 10.0.1.124 -s 10.1.1.78 -u administrator -p mypw`.

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.

- 3 Select a restore point. 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.
- 4 Browse or restore any desired files from the virtual machine.
- 5 When finished browsing or restoring files, enter the command `unmount` and FLR exits.

Troubleshooting VMware Data Recovery

If you have connection or configuration problems with Data Recovery, you can try to resolve these problems by using the suggested troubleshooting solutions.

Table 3-4. Troubleshooting VMware Data Recovery

Problem	Possible Solution
Unable to connect to the backup appliance virtual machine.	There are several possible solutions to this issue, including ensuring that: <ul style="list-style-type: none"> ■ The IPv4 address of the Data Recovery appliance is entered correctly. ■ The client plug-in version matches the backup appliance version. Older client plug-ins may produce errors suggesting incorrectly that the appliance is not powered on. ■ The backup virtual machine is powered on. ■ The ESX/ESXi server hosting the backup appliance can be found on the network. Complications may arise with DNS name resolutions. These issues may be resolved by addressing any DNS name resolution issues or by adding the ESX/ESXi server using an IP address.
Data Recovery fails to complete backups with the error <code>disk full error -1115</code> , 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/ESXi Server and the VMDK is assigned to the appliance. NFS shares cannot be mapped directly to the appliance.
Data Recovery has crashed and the state of the Data Recovery is unknown.	Because the state of the appliance is stored in the deduplication store, it can be restored. Reinstall the Data Recovery appliance to the ESX/ESXi host, and configure the appliance to point to the existing deduplication 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.

Table 3-4. Troubleshooting VMware Data Recovery (Continued)

Problem	Possible Solution
Backup and restore operations are not completing as expected.	<p>An integrity check may have discovered a problem with the integrity of the deduplication store.</p> <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> ■ Maximum of eight backup jobs can run at the same time. ■ Maximum of eight restore jobs can run at the same time. ■ Processor utilization must not exceed 90% to start single backups or 80% to start multiple backups. ■ The datastore where virtual machines are located must have at least 10 GB of space for indexing and processing restore points and 5 GB of storage space available for each virtual machine to be backed up. For example, to simultaneously back up eight virtual machines that reside on one datastore, 50 GB of storage space should be available with 10 GB for indexing and processing and 40 GB for the virtual machines. <p>If any of these limits are exceeded, new jobs do not start.</p>
The Tools status for the Data Recovery backup appliance is listed as unmanaged.	This behavior is expected. The backup appliance is not managed by vCenter Server or other services such as Update Manager. It is not necessary and may not be possible to manage the backup appliance.
Backups fails with error -3960 (cannot quiesce virtual machine)	<p>This may be due to outdated VMware Tools. Ensure the virtual machine to be backed up has the correct version of VMware Tools is installed and up to date. If current tools are not installed, uninstall any existing versions of VMware Tools, and then install the correct version of VMware Tools. This may resolve this issue.</p> <p>If backups continue to fail, try manually creating snapshot of the virtual machine with Snapshot virtual machine's memory unchecked and Quiesce Guest File system checked.</p> <p>For Windows 2003 and later virtual machines, check if system and application event logs for VSS and application writers related messages. Check if ntbackup or Windows Server Backup can be used in the virtual machine to perform backup using VSS in the guest.</p>

Table 3-4. Troubleshooting VMware Data Recovery (Continued)

Problem	Possible Solution
Not all inventory items appear right away after connecting.	If there are a large number of inventory items, some of the items may not appear immediately in the Data Recovery UI. This could occur when the Data Recovery appliance has been powered on within the last few minutes. In this case, wait a few minutes to allow all inventory items to be retrieved before creating or modifying any backup jobs.
Backup jobs do not start as expected.	If the backup appliance was shut down while jobs were in process, jobs may not start again when the appliance is restarted. To avoid this situation, stop all backups using the Data Recovery client, wait for the backups to stop, and then shut down the 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.

Enable Windows 2008 Virtual Machine Application Consistent Quiescing

Windows 2008 virtual machines created on ESX/ESXi 4.0 hosts can be enabled for application consistent quiescing on ESX/ESXi 4.1 and later hosts by enabling the disk UUID attribute.

Procedure

- 1 Start the vSphere Client, and log in to a vCenter Server.
- 2 Select **Virtual Machines and Templates** and click the **Virtual Machines** tab.
- 3 Right-click the Windows 2008 virtual machine for which you are enabling the disk UUID attribute, and select **Power > Power Off**.

The virtual machine powers off.

- 4 Right-click the virtual machine, and click **Edit Settings**.
- 5 Click the **Options** tab, and select the **General** entry in the settings column.
- 6 Click **Configuration Parameters...**

The Configuration Parameters window appears.

- 7 Click **Add Row**.
- 8 In the Name column, enter **disk.EnableUUID**.
- 9 In the Value column, enter **TRUE**.

- 10 Click **OK** and click **Save**.

- 11 Power on the virtual machine.

Application consistent quiescing is available for this virtual machine now that the UUID property has been enabled.

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 reclaiming. 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 drop-down 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 used 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 content of the `datarecovery.ini` file is case-sensitive.

Table 3-5. `datarecovery.ini` Settings

Option	Description	Example	Range	Default
<code>FullIntegrityCheckInterval</code>	The number of days between full integrity checks. If an invalid value is used, it defaults to 7 days.	<code>FullIntegrityCheckInterval=7</code>	1-30.	7
<code>IntegrityCheckInterval</code>	The number of days between integrity checks.	<code>IntegrityCheckInterval=3</code>	1-7.	1
<code>FullIntegrityCheckDay</code>	Can be used to make sure full integrity check only runs on the specified day. 1 = Sunday, 2 = Monday, and so on with 7 = Saturday. If the destination maintenance window does not have any time slot selected on that day, it may cause the destination maintenance to never run and hence should be avoided.	<code>FullIntegrityCheckDay=1</code>		Blank. Allows integrity check to complete on any day.
<code>RetentionPolicyInterval</code>	The number of days between reclaim operations.	<code>RetentionPolicyInterval=3</code>	1-7.	2
<code>SerializeHotadd</code>	Disables parallel SCSI Hot-Add operations, resulting in hot-add operations being completed serially as they were in Data Recovery 1.2 and earlier.	<code>SerializeHotadd=1</code>	0-1.	0

Table 3-5. datarecovery.ini Settings (Continued)

Option	Description	Example	Range	Default
BackupUnusedData	Enables backup of swap partitions in Linux virtual machines and Pagefile.sys in Windows virtual machines. Enabling this attribute causes swap data to be backed up as it was in Data Recovery 1.2 and earlier.	BackupUnusedData=1	0-1.	0
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. If logging is set to higher levels, it may be necessary to increase the MaxLogFiles value so relevant information is not deleted before it can be reviewed.	MaxLogFiles=20		10
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
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

Table 3-5. datarecovery.ini Settings (Continued)

Option	Description	Example	Range	Default
SetBackupSetsLogging	The internal logging level for catalog operations. Higher logging levels are very verbose and should be used only briefly. If higher logging levels are used, consider increasing the MaxLogFiles value.	SetBackupSetsLogging=7	0-7. 7 is most verbose.	3
SetLogging	Overrides logging for all areas.	SetLogging=5	0-7.	None
BackupRetryInterval	The number of minutes to wait before retrying a source for backup.	BackupRetryInterval=20	Any positive integer.	30
DedupeCheckOnRecatalog	Completes an integrity check after a recatalog when set to 1.	DedupeCheckOnRecatalog=1	0-1.	0
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
EnableSVMotionCompatibility	Enables compatibility with Storage vMotion. Set this value to 0 to disable Storage vMotion compatibility, causing Data Recovery to behave as it did in versions 1.2.1 and earlier.	EnableSVMotionCompatibility=1	0-1.	1
MaxBackupRestoreTasks	The maximum number of simultaneous backup and restores.	MaxBackupRestoreTasks=4	1-8.	8
ConnectionAcceptMode	Determines the types of connections the backup appliance accepts from the vSphere Client plugin. 1 requires SSL connections. 2 requires plaintext connections. 3 attempts SSL connections, but allows plaintext if SSL is not supported by the client. If your environment includes older vSphere Client plug-ins, either upgrade those plug-ins or permit plaintext. Requiring SSL connections with older client plug-ins result in connection attempt failures.	ConnectionAcceptMode=2	1-3.	1

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 Click the IP address text field.
- 3 Enter the keystroke series Ctrl-Alt-g-g.
The client connection logs are displayed.

Index

A

- adding
 - network share **28**
 - storage **18**

B

- backup
 - manual **28**
 - process **7**
 - scaling **7**
- backup appliance
 - configuring **24**
 - connect to vcenter server **25**
 - installing **17**
 - power on **23**
- backup job
 - creating **26, 27**
 - options **26**
- backup job wizard, using **27**
- backups, suspend **29**
- bring to compliance **28**

C

- client, installing **17**
- client connection logs, viewing **45**
- configuring
 - backup appliance **24**
 - data recovery **21**
- creating, backup job **26**

D

- damaged restore points
 - remove **41**
 - understanding **41**
- data recovery
 - configuring **21**
 - prerequisites **13**
 - scaling **13**
- data recovery logs, using **45**
- datarecover.ini, reference **42**
- datarecovery.ini
 - modify backup appliance behavior **41**
 - understanding **41**
- deduplication
 - best practices **10**
 - scaling **10**
- disk, extend **19**

E

- email, reporting **30**
- extend, disk **19**

F

- file level restore, See flr
- firewalls **17**
- flr, understanding **33**
- flr,advanced mode,linux **37**

G

- getting started wizard, using **25**

I

- install flr
 - linux **36**
 - windows **35**
- installing
 - backup appliance **17**
 - client **17**
 - data recovery **13**
- integrity check **10**
- introducing, data recovery **7**

L

- licensing **7**
- logs, viewing **45**

M

- maintenance operations, schedule **29**

N

- network share, adding **28**

P

- preface **5**

R

- recatalog **10**
- reclaim **10**
- reporting, email **30**
- restore files, windows **35**
- restore files using flr
 - linux **37**
 - windows **35**
- restore points, mark for removal or locking **29**
- restore rehearsal **31, 32**

restoring, virtual machines **31, 32**

S

scaling

 backup **7**

 data recovery **13**

 deduplication **10**

schedule, maintenance operations **29**

storage, adding **18**

supported storage **7**

suspend, backups **29**

T

troubleshooting **38**

U

understanding, fir **33**

user interface, understanding **21**

using, getting started wizard **25**

V

verbose logs, view **45**

virtual machines, restoring **31, 32**

Volume Shadow Copy Service, *See also* VSS

volumes, formatting **28**

VSS

 benefits **8**

 support **8**

 understanding **8**

W

windows 2008 virtual machine application
 consistent quiescing, enabling **40**