

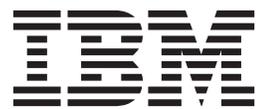
IBM Tivoli Storage Manager for Virtual Environments
Version 6.3

*Data Protection for VMware Installation
and User's Guide*



IBM Tivoli Storage Manager for Virtual Environments
Version 6.3

*Data Protection for VMware Installation
and User's Guide*



Note:

Before using this information and the product it supports, read the information in "Notices" on page 109.

This edition applies to Version 6.3 of Tivoli Storage Manager for Virtual Environments (product number 5725-A44) and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC27-2898-00.

Changes since the previous edition are marked with a vertical bar (|) in the left margin. Ensure that you are using the correct edition for the level of the product.

© **Copyright IBM Corporation 2011.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	v
Who should read this guide	v
Publications	v
Tivoli Storage Manager publications	v
Tivoli Storage FlashCopy Manager publications	viii
Support information	viii
Getting technical training	viii
Searching knowledge bases	ix
Contacting IBM Software Support	xi

What's new for Tivoli Storage Manager for Virtual Environments 6.3 xiii

Chapter 1. Tivoli Storage Manager for Virtual Environments overview 1

System components	2
-----------------------------	---

Chapter 2. Planning. 5

Supported operating systems	5
Supported operating systems for the Recovery Agent command-line interface	6
Supported operating systems for Data Protection for VMware Recovery Agent	7
Supported operating systems for Data Protection for VMware vCenter plug-in	10
Hardware requirements	13
Software requirements and prerequisites.	13
Virtual environment configurations	17

Chapter 3. Installing, upgrading, and uninstalling Data Protection for VMware 23

Preparing for installation	23
Installing Data Protection for VMware on Windows	24
Installing Data Protection for VMware on Linux	26
Performing a clean installation of Data Protection for VMware on Linux	28
Installing language packs	28
Installing a language pack on Windows	28
Installing a language pack on Linux	29
Installing Data Protection for VMware in silent mode	29
Installing Data Protection for VMware on a Windows 32-bit system in silent mode	30
Installing Data Protection for VMware on a Windows 64-bit system in silent mode	30
Installing Data Protection for VMware on a Linux system in silent mode	31
Upgrading Data Protection for VMware	32
Upgrading Data Protection for VMware on a Windows 32-bit system in silent mode	33
Upgrading Data Protection for VMware on a Windows 64-bit system in silent mode	34

Upgrading Data Protection for VMware on a Linux system in silent mode.	34
Uninstalling Data Protection for VMware	35
Uninstalling Data Protection for VMware for a Windows 32-bit system in silent mode	35
Uninstalling Data Protection for VMware for Windows 64-bit system in silent mode	35
Uninstalling Data Protection for VMware a Linux system	36

Chapter 4. Starting and configuring Data Protection for VMware. 39

Starting and running services for Data Protection for VMware	39
Configuration and log files	40

Chapter 5. Using the Data Protection for VMware vCenter plug-in. 41

Configuring the Data Protection for VMware vCenter plug-in	41
Register Tivoli Storage Manager nodes	41
Setting up the Tivoli Storage Manager Backup-Archive Client on the vStorage Backup Server for each datamover node	44
Configuring the Data Protection for VMware command-line interface	46
Start the Data Protection for VMware vCenter plug-in	48
Backing up virtual machine data to Tivoli Storage Manager	49
Back up options	50
Restoring a single virtual machine from Tivoli Storage Manager	51
Restoring multiple virtual machines from Tivoli Storage Manager	52
Reporting	53
Troubleshooting	54

Chapter 6. Using the Data Protection for VMware command-line interface . . . 57

Backup	57
Restore	59
Inquire_config	61
Inquire_detail.	62
Set_password.	63
Profile parameters	64

Chapter 7. Using the Data Protection for VMware Recovery Agent 69

Backing up VMware virtual machine data	69
Scenario: Backing up your virtual machines	69
Preparing the environment for VMware backup processing.	70
Running full virtual machine backups	73

Running incremental virtual machine backups using the VStore API	74
Restoring full VM backups	74
Backup VM	75
Restoring virtual machine data	77
Mounting snapshots	77
Tape configuration guidelines	81
Configuring systems for iSCSI mount.	84
Restoring files and instant restore of volumes	85

Chapter 8. Recovery Agent command-line interface 99

Starting the Recovery Agent command-line interface	99
Recovery Agent command-line interface overview	99
mount.	101
set_connection	104

help	104
Recovery Agent command-line interface return codes	104

Appendix. Accessibility features for Tivoli Storage Manager for Virtual Environments 107

Notices	109
Trademarks	111

Glossary 113

Index 135

Preface

IBM® Tivoli® Storage Manager for Virtual Environments provides off-host single source block incremental backup and file recovery and instant restore from a full-VM backup for Windows and Linux guest machines. In addition, Data Protection for VMware allows the backup-archive client to perform block level incremental backups.

This publication describes how to install, configure, and use Data Protection for VMware.

Who should read this guide

This guide is intended for anyone who wants to use Tivoli Storage Manager for Virtual Environments.

Publications

Publications for the IBM Tivoli Storage Manager family of products are available online. The IBM Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search all publications, go to the Tivoli Storage Manager information center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>.

You can download PDF versions of publications from the Tivoli Storage Manager information center or from the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>.

You can also order some related publications from the IBM Publications Center website. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	GC23-9781
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC23-9769

Table 1. Tivoli Storage Manager server publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC23-9775
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	GC23-9782
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i>	SC23-9770
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC23-9776
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	GC23-9783
<i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i>	SC23-9771
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC23-9777
<i>IBM Tivoli Storage Manager for Oracle Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference</i>	SC23-9778
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	GC23-9785
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC23-9773
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC23-9779
<i>IBM Tivoli Storage Manager for z/OS Media Installation and User's Guide</i>	SC27-4018
<i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i>	GC27-4017
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack®</i>	SC27-2828

Table 2. Tivoli Storage Manager storage agent publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i>	SC23-9797
<i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i>	SC23-9798
<i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i>	SC23-9799
<i>IBM Tivoli Storage Manager for SAN for Oracle Solaris Storage Agent User's Guide</i>	SC23-9800
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 3. Tivoli Storage Manager client publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i>	SC23-9791
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i>	SC23-9795

Table 4. Tivoli Storage Manager data protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	GC27-4010
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC27-4019
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC27-4020
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	GC27-4009
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino® UNIX and Linux Installation and User's Guide</i>	SC27-4021
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC27-4022
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i>	SC33-6341
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i>	SC33-6340
<i>IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide</i>	SC27-2898
<i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i>	N/A

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i>	GC27-4011
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i>	GC27-4012
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i>	SC27-4014
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i>	SC27-4015
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages</i>	SC27-4016

Note: You can find information about IBM System Storage® Archive Manager at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/c_complydataretention_ovr.html.

Tivoli Storage FlashCopy Manager publications

The following table lists the publications that make up the Tivoli Storage FlashCopy Manager library.

Table 6. Tivoli Storage FlashCopy Manager publications

Publication title	Order number
<i>IBM Tivoli Storage FlashCopy Manager for UNIX and Linux Installation and User's Guide</i>	SC27-4005
<i>IBM Tivoli Storage FlashCopy Manager for Windows Installation and User's Guide</i>	SC27-4006
<i>IBM Tivoli Storage FlashCopy Manager for VMware Installation and User's Guide</i>	SC27-4007
<i>IBM Tivoli Storage FlashCopy Manager Messages</i>	GC27-4008

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

If you have a problem with your Tivoli Storage Manager family product, there are several knowledge bases that you can search.

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

If you cannot find an answer to your question in the IBM Tivoli Storage Manager information center, search the Internet for the information that might help you resolve your problem.

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to the following wikis:

Tivoli Storage Manager wiki

<http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>

Tivoli Storage FlashCopy Manager wiki

[https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage FlashCopy Manager](https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager)

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.
6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xii.

Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?

- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

What's new for Tivoli Storage Manager for Virtual Environments 6.3

The Tivoli Storage Manager for Virtual Environments product compliments other Tivoli Storage Manager products. Data Protection for VMware version 6.3 contains many new features and changes.

Linux Windows **Data Protection for VMware vCenter plug-in**

This component is a new graphical user interface (GUI) that integrates with the VMware vSphere Client. The Data Protection for VMware vCenter plug-in is accessed as a vCenter Server extension in the Solutions and Applications panel of your vCenter Server System. Use the Data Protection for VMware vCenter plug-in as the primary interface from which to complete these tasks:

- Initiate a backup of your virtual machines to a Tivoli Storage Manager server (or schedule a backup for a later time).
- Initiate a full recovery of your virtual machines from a Tivoli Storage Manager server.
- Issue reports concerning backup, restore, and configuration activity.

Chapter 5, “Using the Data Protection for VMware vCenter plug-in,” on page 41 provides requisite information, including detailed configuration and user tasks.

Linux Windows **Data Protection for VMware command-line interface**

This component (Data Protection for VMware CLI) is a new full-function command-line interface that is installed with the Data Protection for VMware vCenter plug-in. You can use the Data Protection for VMware CLI to complete these tasks:

- Initiate a backup of your virtual machines to a Tivoli Storage Manager server (or schedule a backup for a later time).
- Initiate a full recovery of your virtual machines, virtual machine files, or Virtual Machine Disks (VMDKs) from a Tivoli Storage Manager server.
- View configuration information about the backup database and environment.

Although the Data Protection for VMware vCenter plug-in is the primary task interface, the Data Protection for VMware CLI provides a useful secondary interface. For example, the Data Protection for VMware CLI can be used to implement a scheduling mechanism different from the one implemented by the Data Protection for VMware vCenter plug-in. Also, the Data Protection for VMware CLI is useful when evaluating automation results with scripts. Chapter 6, “Using the Data Protection for VMware command-line interface,” on page 57 provides detailed information regarding available commands.

Linux Solaris **Additional Linux and Solaris support for Data Protection for VMware Recovery Agent**

Data Protection for VMware Recovery Agent mount and restore functions are now supported on Red Hat Enterprise Linux 6.x Servers and SUSE Linux Enterprise Server 11.

Chapter 1. Tivoli Storage Manager for Virtual Environments overview

IBM Tivoli Storage Manager for Virtual Environments (referred to as Data Protection for VMware) protects virtual machines by offloading backup workloads to a centralized server and providing multiple restore capabilities.

Full VM restore

Restore a full or incremental VM backup. The entire VM is restored to the state it existed in when originally backed up.

Full VM backups

Back up an entire virtual machine in your vCenter to Tivoli Storage Manager storage as a single object (image). This is an entire VM image snapshot, which is a single snapshot that contains all of the VM disks. All data is backed up at the disk block level. The data can then be restored to a disk, or mounted as a virtual volume for an individual file restore. These backups are managed and retained according to storage policies set up by the Tivoli Storage Manager administrator.

Incremental VM backups

Back up only the virtual machine data that has changed since the last full backup completed. All data is backed up at the disk block level. When a full backup has not been issued for the virtual machine, Data Protection for VMware issues a full backup by default.

You can also use Data Protection for VMware to interface with the Windows Backup-Archive Client in order to use the Change Block Tracking capabilities provided by the vStorage APIs for Data Protection (VADP).

File level restore

File level restore can be performed in-guest or off-host, and on supported Windows and Linux machines. Mount the volumes of the VM backup as a virtual volume. Then, copy the files that you want to restore using tools such as Windows Explorer or system commands. File restore is done from a Backup-Archive Client full or incremental virtual machine backup. The recovery point represented by either a full or incremental backup can be mounted. Although the mounted volume is read only by default, write permissions are also available.

Instant restore

With instant restore, you can restore the content of a single volume from a snapshot. This restore uses the snapshot data generated by the Backup-Archive Client. Instant restore can be done from a full or incremental virtual machine backup. You can use the volume immediately, while the restore process continues in the background.

This component (Data Protection for VMware CLI) is a new full-function command-line interface that is installed with the Data Protection for VMware vCenter plug-in. You can use the Data Protection for VMware CLI to complete these tasks:

- Initiate a backup of your virtual machines to a Tivoli Storage Manager server (or schedule a backup for a later time).
- Initiate a full recovery of your virtual machines, virtual machine files, or Virtual Machine Disks (VMDKs) from a Tivoli Storage Manager server.
- View configuration information about the backup database and environment.

Although the Data Protection for VMware vCenter plug-in is the primary task interface, the Data Protection for VMware CLI provides a useful secondary interface. For example, the Data Protection for VMware CLI can be used to implement a scheduling mechanism different from the one implemented by the Data Protection for VMware vCenter plug-in. Also, the Data Protection for VMware CLI is useful when evaluating automation results with scripts.

For detailed information regarding available commands, see Chapter 6, "Using the Data Protection for VMware command-line interface," on page 57.

Data Protection for VMware Recovery Agent

This service enables the mounting of any snapshot volume from the Tivoli Storage Manager server. You can view the snapshot locally, with read-only access, on the client system, or use an iSCSI protocol to access the snapshot from a remote computer. In addition, the Data Protection for VMware Recovery Agent provides the instant restore function. A volume used in instant restore processing remains available while the restore process proceeds in the background.

For detailed information regarding tasks, scenarios, and guidelines, see Chapter 7, "Using the Data Protection for VMware Recovery Agent," on page 69.

Data Protection for VMware Recovery Agent command-line interface

You can use this component (Recovery Agent CLI), which is installed on Windows only, to perform the following tasks from a remote machine:

- Gather information about available restorable data, including lists of:
 - Backed-up virtual machines
 - Snapshots available for a backed-up machine
 - Partitions available in a specific snapshot
- Mount a snapshot as a virtual device.
- Get a list of virtual devices.
- Remove a virtual device.

For detailed information regarding TDPVMwareShell.exe commands, parameters, and return codes, see Chapter 8, "Recovery Agent command-line interface," on page 99.

Note: In Tivoli Storage Manager for Virtual Environments version 6.2, this Data Protection for VMware Recovery Agent command-line interface was named the "command-line interface".

Chapter 2. Planning

Before you install Data Protection for VMware, verify that your system is running a supported operating system, and that you meet all hardware and software requirements.

Data Protection for VMware supports any disk configuration that is supported by the hardware and operating system. The disk configuration includes multipath device drivers.

Supported operating systems

To implement Data Protection for VMware components, your site must have the appropriate operating system and environment, hardware, and software.

Table 7. Available Data Protection for VMware components by operating system

Operating system	Recovery Agent command-line interface	Data Protection for VMware Recovery Agent	Data Protection for VMware vCenter plug-in ¹
Microsoft Windows 2003	√	√	√
Microsoft Windows 2003 (64-bit Edition)	√	√	√
Microsoft Windows 2008	√	√	√
Microsoft Windows 2008 R2 (or later)	√	√	√
Microsoft Windows Vista	√	√	
Microsoft Windows XP	√	√	
Microsoft Windows 7	√	√	√
Red Hat Enterprise Linux 5.2, 5.3, 5.4, 5.5, 5.6 servers		√	√ ²
Red Hat Enterprise Linux 6.x servers		√	√ ²
SUSE Linux Enterprise Server 10		√	√ ²
SUSE Linux Enterprise Server 11		√	√ ²

See the following sections for detailed information regarding operating system support:

- “Supported operating systems for the Recovery Agent command-line interface” on page 6
- “Supported operating systems for Data Protection for VMware Recovery Agent” on page 7
- “Supported operating systems for Data Protection for VMware vCenter plug-in” on page 10

Note:

1. Includes the Data Protection for VMware command-line interface.
2. Available on 64-bit operating systems only.

Supported operating systems for the Recovery Agent command-line interface

Ensure that you are installing the Recovery Agent CLI on a supported operating system.

The following table provides details about operating systems that are supported for the Recovery Agent CLI.

Table 8. Operating systems for the command line

Operating system and supported release	Support details
Microsoft Windows 2003, Service Pack 1 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Storage Server • Storage R2 Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
Microsoft Windows 2003 64-bit Edition	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) and IA64 (Intel Itanium) instruction set architecture • Supports 64-bit processors
Microsoft Windows 2008, Service Pack 1 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Datacenter Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit), x64 (AMD64 and EM64T), and IA64 (Intel Itanium) instruction set architecture • Supports 32-bit and 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.
Microsoft Windows 2008, R2 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Datacenter Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture • Supports 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.
Microsoft Windows 7 for all editions	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors

Table 8. Operating systems for the command line (continued)

Operating system and supported release	Support details
Microsoft Windows Vista, Service Pack 1 or later: <ul style="list-style-type: none"> • Starter • Home Basic • Home Premium • Business • Enterprise • Ultimate 	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
Microsoft Windows XP Professional Edition, Service Pack 2 or later	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit processors

Supported operating systems for Data Protection for VMware Recovery Agent

Ensure that you are installing Data Protection for VMware Recovery Agent on a supported operating system.

The following table provides details about operating systems that are supported for Data Protection for VMware Recovery Agent.

Table 9. Operating systems for Data Protection for VMware Recovery Agent

Operating system and supported release	Support details
Microsoft Windows 2003, Service Pack 1 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Storage Server • Storage R2 Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
Microsoft Windows 2003 64-bit Edition	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture • Supports 64-bit processors
Microsoft Windows 2008, Service Pack 1 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit), x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.

Table 9. Operating systems for Data Protection for VMware Recovery Agent (continued)

Operating system and supported release	Support details
<p>Microsoft Windows 2008, R2 or later for the following servers:</p> <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture. • Supports 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.
<p>Microsoft Windows 7 for all editions</p>	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
<p>Microsoft Windows Vista, Service Pack 1 or later:</p> <ul style="list-style-type: none"> • Starter • Home Basic • Home Premium • Business • Enterprise • Ultimate 	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
<p>Microsoft Windows XP, Professional Edition, Service Pack 2 or later</p>	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit processors
<p>Red Hat Enterprise Linux 5.2, 5.3, 5.4, 5.5, 5.6 servers</p>	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit and 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – RedHat-i386: 2.6.18-92.e15.i686 and 2.6.18-92.e15.i686 PAE – RedHat-x86_64: 2.6.18-92.e15-x86_64 • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux package <code>iscsi-initiator-utils-6.2.0.868-0.7.el5</code> • <code>lsscsi</code> command • <code>iscsiadm</code> utility • Secure Shell (SSH) client for Linux

Table 9. Operating systems for Data Protection for VMware Recovery Agent (continued)

Operating system and supported release	Support details
Red Hat Enterprise Linux 6.x servers	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit and 64-bit processors • Requires the following minimum packages and their dependencies: code>compat-libstdc++-33-3.2.3-69 code>compat-db-4.6.21-15 code>libXp-1.0.0-15.1 code>libXmu-1.0.5-1 code>libXtst-1.0.99.2-3 code>pam-1.1.1-4 code>libXft-2.1.13-4.1 code>gtk2-2.18.9-4 code>gtk2-engines-2.18.4-5 Installations on 64-bit processors require the 32-bit versions of these packages. • The following kernels are supported: <ul style="list-style-type: none"> – RedHat-i386: 2.6.18-92.e15.i686 and 2.6.18-92.e15.i686 PAE – RedHat-x86_64: 2.6.18-92.el5-x86_64 • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux package <code>iscsi-initiator-utils-6.2.0.868-0.7.el5</code> • <code>lsscsi</code> command version 0.23 (or later) • <code>iscsiadm</code> utility • Secure Shell (SSH) client for Linux
SUSE Linux Enterprise Server 10, Service Pack 2	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit and 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – SUSE-i386: 2.6.16.60-0.21default, 2.6.16.60-0.21smp, and 2.6.16.60-0.21bigsm – SUSE-x86_64: 2.6.16.60-0.21default and 2.6.16.60-0.21smp For all kernel versions, auto mount is not supported. • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux • <code>lsscsi</code> command • <code>iscsiadm</code> utility • Secure Shell (SSH) client for Linux

Table 9. Operating systems for Data Protection for VMware Recovery Agent (continued)

Operating system and supported release	Support details
SUSE Linux Enterprise Server 11	<ul style="list-style-type: none"> • Supports the x86 (32-bit) instruction set architecture • Supports 32-bit and 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – SUSE-i386: 2.6.16.60-0.21default, 2.6.16.60-0.21smp, and 2.6.16.60-0.21bigsmpt – SUSE-x86_64: 2.6.16.60-0.21default and 2.6.16.60-0.21smp <p>For all kernel versions, auto mount is not supported.</p> <ul style="list-style-type: none"> • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux • lsscsi command • iscsiadm utility • Secure Shell (SSH) client for Linux

Note: Windows Support is not provided for applications that use SCSI Pass Through Interface (SPTI) or SCSI Pass Through Direct (SPTD) for performing read and write operations. You cannot use instant restore while applications that use SPTI or SPTD are running. If you try to use instant restore while applications that use SPTI or SPTD are running, it might appear that the instant restore was completed, but the data might be corrupted.

Supported operating systems for Data Protection for VMware vCenter plug-in

Ensure that you are installing Data Protection for VMware vCenter plug-in on a supported operating system.

The following table provides details about operating systems that are supported for the Data Protection for VMware vCenter plug-in.

Table 10. Operating systems for Data Protection for VMware vCenter plug-in

Operating system and supported release	Support details
Microsoft Windows Server 2003, Service Pack 1 or later for the following servers: <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Storage Server • Storage R2 Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
Microsoft Windows Server 2003 64-bit Edition	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture • Supports 64-bit processors

Table 10. Operating systems for Data Protection for VMware vCenter plug-in (continued)

Operating system and supported release	Support details
<p>Microsoft Windows Server 2008, Service Pack 1 or later for the following servers:</p> <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x86 (32-bit), x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.
<p>Microsoft Windows Server 2008, R2 or later for the following servers:</p> <ul style="list-style-type: none"> • Standard Server • Enterprise Server • Web Server • Storage Server • Small Business Server • Essential Business Server 	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture. • Supports 64-bit processors • If you use Active Directory with Microsoft Windows 2008, see the Microsoft Knowledge Base article 970770 online at http://support.microsoft.com/default.aspx?scid=kb;EN-US;970770 . Download the hotfix associated with this knowledge base article.
<p>Microsoft Windows 7 for all editions</p>	<ul style="list-style-type: none"> • Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture • Supports 32-bit and 64-bit processors
<p>Red Hat Enterprise Linux 5.2, 5.3, 5.4, 5.5, 5.6 servers</p>	<ul style="list-style-type: none"> • Supports the x86_64 instruction set architecture • Supports 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – RedHat-i386: 2.6.18-92.e15.i686 and 2.6.18-92.e15.i686 PAE – RedHat-x86_64: 2.6.18-92.el5-x86_64 • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux package iscsi-initiator-utils-6.2.0.868-0.7.el5 • lsscsi command • iscsiadm utility • Secure Shell (SSH) client for Linux

Table 10. Operating systems for Data Protection for VMware vCenter plug-in (continued)

Operating system and supported release	Support details
Red Hat Enterprise Linux 6.x servers	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture • Supports 64-bit processors • Requires the following minimum packages and their dependencies: compat-libstdc++-33-3.2.3-69 compat-db-4.6.21-15 libXp-1.0.0-15.1 libXmu-1.0.5-1 libXtst-1.0.99.2-3 pam-1.1.1-4 libXft-2.1.13-4.1 gtk2-2.18.9-4 gtk2-engines-2.18.4-5 • Installations on 64-bit processors require the 32-bit versions of these packages. • The following kernels are supported: <ul style="list-style-type: none"> – RedHat-i386: 2.6.18-92.e15.i686 and 2.6.18-92.e15.i686 PAE • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux package iscsi-initiator-utils-6.2.0.868-0.7.el5 • lsscsi command version 0.23 (or later) • iscsiadm utility • Secure Shell (SSH) client for Linux
SUSE Linux Enterprise Server 10, Service Pack 2	<ul style="list-style-type: none"> • Supports the x64 (AMD64 and EM64T) instruction set architecture • Supports 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – SUSE-i386: 2.6.16.60-0.21default, 2.6.16.60-0.21smp, and 2.6.16.60-0.21bigsm – SUSE-x86_64: 2.6.16.60-0.21default and 2.6.16.60-0.21smp • For all kernel versions, auto mount is not supported. • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux • lsscsi command • iscsiadm utility • Secure Shell (SSH) client for Linux

Table 10. Operating systems for Data Protection for VMware vCenter plug-in (continued)

Operating system and supported release	Support details
SUSE Linux Enterprise Server 11	<ul style="list-style-type: none"> • Supports the x86_64 (AMD64 and EM64T) and IA64 (Intel Itanium) instruction set architecture • Supports 64-bit processors • The following kernels are supported: <ul style="list-style-type: none"> – SUSE-i386: 2.6.16.60-0.21default, 2.6.16.60-0.21smp, and 2.6.16.60-0.21bigmp – SUSE-x86_64: 2.6.16.60-0.21default and 2.6.16.60-0.21smp For all kernel versions, auto mount is not supported. • Perl version 5 on Linux systems • mdadm tool for managing Linux Software RAID arrays • iSCSI Initiator for Linux • lsscsi command • iscsiadm utility • Secure Shell (SSH) client for Linux

Hardware requirements

Hardware requirements vary and depend on the following items:

- Number of protected servers
- Number of protected volumes
- Data set sizes
- LAN and SAN connectivity

Note: Data Protection for VMware does not support operations in a LAN-free environment.

The following table describes the hardware requirements that are needed to install Data Protection for VMware.

Table 11. Hardware requirements for Data Protection for VMware.

Component	Minimal requirement	Preferred
System	3 GHz Dual Intel Pentium D processor or compatible	Not applicable
Memory	2 GB RAM, 2 GB virtual address space	Not applicable
Available hard disk	200 MB for 'Documents and Settings' folder	2 GB
NIC Card	1 NIC - 100 Mbps	1 NIC - 1 Gbps

Software requirements and prerequisites

Before installing Data Protection for VMware Version 6.3, some applications, utilities, and components must be installed or available.

Data Protection for VMware vCenter plug-in

The VMware vSphere client uses the installed Internet Explorer as the default browser to run the Data Protection for VMware vCenter plug-in. The supported browser for the Data Protection for VMware vCenter plug-in is Internet Explorer 8.

Data Protection for VMware Recovery Agent

Data Protection for VMware Recovery Agent uses an internal Tivoli Storage Manager protocol to connect to the server. Port 1500 is the default port that Tivoli Storage Manager uses for Data Protection for VMware Recovery Agent to work. You can customize the port.

Users must be logged in locally in order to run operations from the Data Protection for VMware Recovery Agent GUI. When issuing commands using the Data Protection for VMware Recovery Agent command-line interface, users can be logged in locally or logged in on a remote machine. Data Protection for VMware Recovery Agent can be used when it is accessed through a remote desktop when connecting in console mode, by using the administrator switch.

Using the command line to drive Data Protection for VMware Recovery Agent on a Linux machine

To use the command line client from a system running a supported Linux operating system, complete the following prerequisite configuration tasks:

1. On the Windows system where you have installed or plan to install the command line client, install Cygwin 1.5.25 or later. When you install Cygwin, include the OpenSSH package. To manually install Cygwin, complete the following steps:
 - a. Log on to the Windows server using an account with administrator privileges.
 - b. Go to the following web site and install Cygwin 1.5.25 or later:
<http://www.cygwin.com>
 - c. When completing the installation wizard for Cygwin, there is a **Select Package** page. On this page, clear the **Hide obsolete and administrative packages** check box.
 - d. During the installation process for Cygwin, select the following Cygwin packages:

Table 12. Cygwin packages

Category	Package
Net	All default packages. In addition, select the following packages: <ul style="list-style-type: none">• openssh (contains ssh.exe)• openssl (contains ssl.exe)• rsync• tcp_wrappers

- e. After finishing the Cygwin installation wizard, add the `Cygwin\bin` directory to the Microsoft Windows `%PATH%` environment variable. The directory must be the first one in the `%PATH%` environment variable.

Remember: Restart the system so the variable update can take effect.

2. On the system where you have installed or plan to install the command line client, test the Cygwin installation.

Remember: Before using Cygwin, review the Cygwin documentation for any issues that might affect your environment.

To test the Cygwin installation, from the Microsoft Windows Start menu, select **Programs > Cygnus Solutions > Cygwin Bash Shell**. A command prompt window should be displayed. This window is a bash shell.

3. On the system where you have installed or plan to install the command line client with Cygwin, install the SSH daemon service. To install the SSH daemon service, complete the following steps:
 - a. Enter the following commands to give read, write, and file owner permissions to the `/etc/passwd` and `/etc/group` files:

```
chmod r+u+w /etc/passwd
chmod r+u+w /etc/group
chmod 664 /var/log/sshd.log
```
 - b. Enter the following command to give read access to the `/var` directory:

```
chmod 755 /var
```
 - c. From the Cygwin command prompt window, run the following command to create the SSH daemon service:

```
ssh-host-config
```
 - d. When a query about whether privilege separation should be used is displayed in the command prompt window, enter *no*.
 - e. When a query about whether a new local account named `sshd` should be created is displayed in the command prompt window, enter *yes*.
 - f. When a query about whether `sshd` should be installed as a service is posted in the command prompt window, enter *yes*.
 - g. When a query asks you to enter the value of **CYGWIN** for the daemon, enter the following text: *ntsec tty*
 - h. When a query asks if you want to use a different name, enter *no*.
 - i. When a query asks if you want to create a new privileged user account named `cyg_server`, enter *yes*.
 - j. When a query asks you to enter a password, enter a password. You are asked to reenter the password to confirm the entry. The host configuration is complete. A status message is displayed.
 - k. At the prompt, enter the following command:

```
set CYGWIN 'ntsec tty'
```

Also, add CYGWIN as a Microsoft Windows environment variable with the value *ntsec tty*.

4. Configure the authentication key files by logging on to the Linux system where Data Protection for VMware is installed and by completing these tasks:
 - a. Issue this command and press **Enter** at all prompt questions:

```
ssh-keygen -t dsa
```
 - b. Issue these commands:

```
cd .ssh
scp id_dsa.pub Administrator@windows_machine:/home/Administrator
```
 - c. Issue these commands from the Cygwin shell on the Windows server:

```

mkdir .ssh
chmod 700 .ssh
cd .ssh
touch authorized_keys
cat ../id_dsa.pub >> authorized_keys
rm ../id_dsa.pub

```

- d. Configure the SSH server to use the authentication files by editing the SSH service configuration file `c:\cygwin\etc\sshd_config`. Open this file and unmark these entries:

```

Protocol 2
HostKey /etc/ssh_host_dsa_key
RSAAuthentication yes
PubkeyAuthentication yes
AuthorizedKeysFile

```

Update the `AuthorizedKeysFile` value to specify `/home/Administrator/.ssh/authorized_keys`.

- e. Issue these commands from the Cygwin shell on the Windows server to restart the `sshd` service:

```

net stop sshd
net start sshd

```

- f. Verify that the Linux system can communicate with the Windows server system by issuing this command (from the Linux system):

```
ssh Administrator@windows_machine
```

SSH attempts to update the `known_hosts` file for each host name convention specified. For example, although each of these commands identify the same Windows Server, SSH attempts to add an entry to the `known_hosts` file for each host name:

```
ssh Administrator@windows_machine
ssh Administrator@windows_machine.xyz.com

```

To prevent possible timeout errors due to authentication failure, implement one (or both) of these recommendations:

- Consistently use the same host name convention when accessing the Windows Server.
- Update the `known_hosts` file with all host name conventions associated with the Windows Server.

Important: You must create authentication key files for each new client system. Therefore, complete Steps 4a through 4f for each client system.

5. Permit any host to connect using SSH to the server by editing the following file: `C:\cygwin\etc\hosts.allow`

The following line must immediately precede the `ALL : PARANOID : deny` line:

```
sshd: ALL
```

6. Log on to the command line (without a password).
 7. In the command prompt window, enter the following command:

```
TDPVMwareShell.exe -c command type tag parameter
```

In addition to the Cygwin and SSH daemon service, the GNU C libraries, Version 2.3.3-98.38 or later are required.

Virtual environment configurations

Data Protection for VMware provides a variety of configurations for performing file-level restore, instant restore, and disk / block device exposure.

Off-host file-level restore for Windows and for Linux

These configurations do not require Data Protection for VMware Recovery Agent to be installed in each virtual machine guest. Instead, an off-host Windows or Linux instance is responsible for file-level restore of multiple virtual machines. With this configuration, the mount process exposes a virtual volume from a selected disk partition.

The Tivoli Storage Manager administrator must register a node that is associated with the Data Protection for VMware Recovery Agent. This Tivoli Storage Manager node name that is associated with the Data Protection for VMware Recovery Agent instance requires access to all virtual machines. Issue the following command from the backup-archive client node that owns the virtual machines:

```
set access backup * mountnodename
```

When a snapshot has been mounted to the off-host server, it can then be exported so that a virtual machine guest user can access to the files. This centralized restore process is typically initiated by a VMware administrator, by a Tivoli Storage Manager administrator, or by help desk personnel.

For these configurations, ensure that you compare the specific virtual machine guest operating system requirements with the supported levels of Data Protection for VMware Recovery Agent. If a specific operating system is not supported, determine if the off-host disk / block device exposure configuration could be used. See Figure 8 on page 21.

The data paths for off-host file restores are illustrated in Figure 2 and Figure 3 on page 18

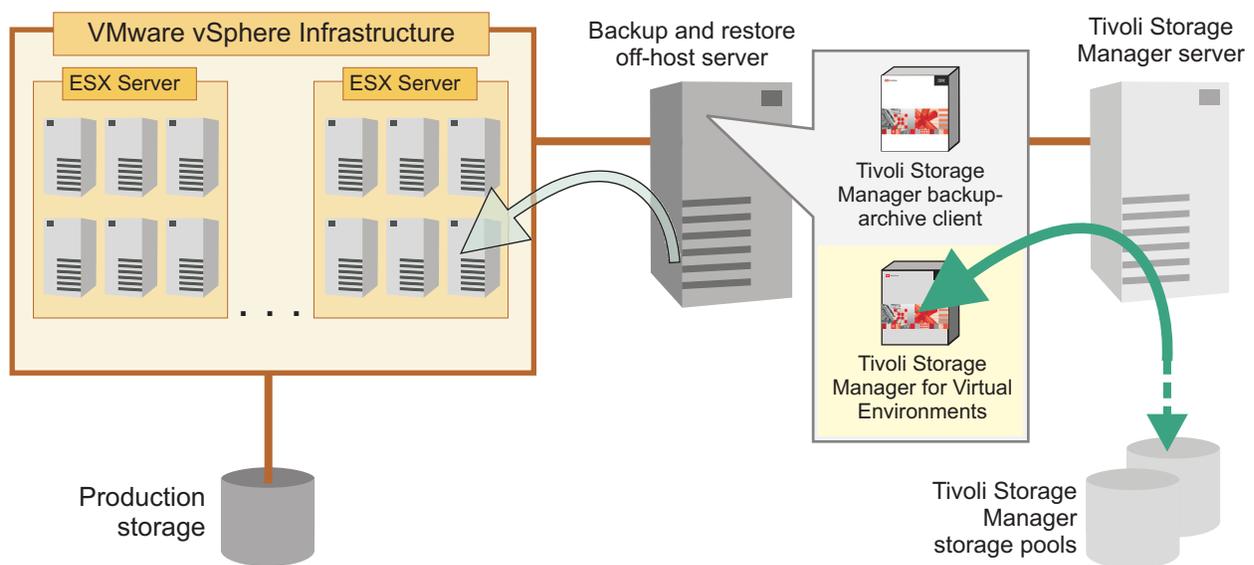


Figure 2. Off-host file-level restore for Windows

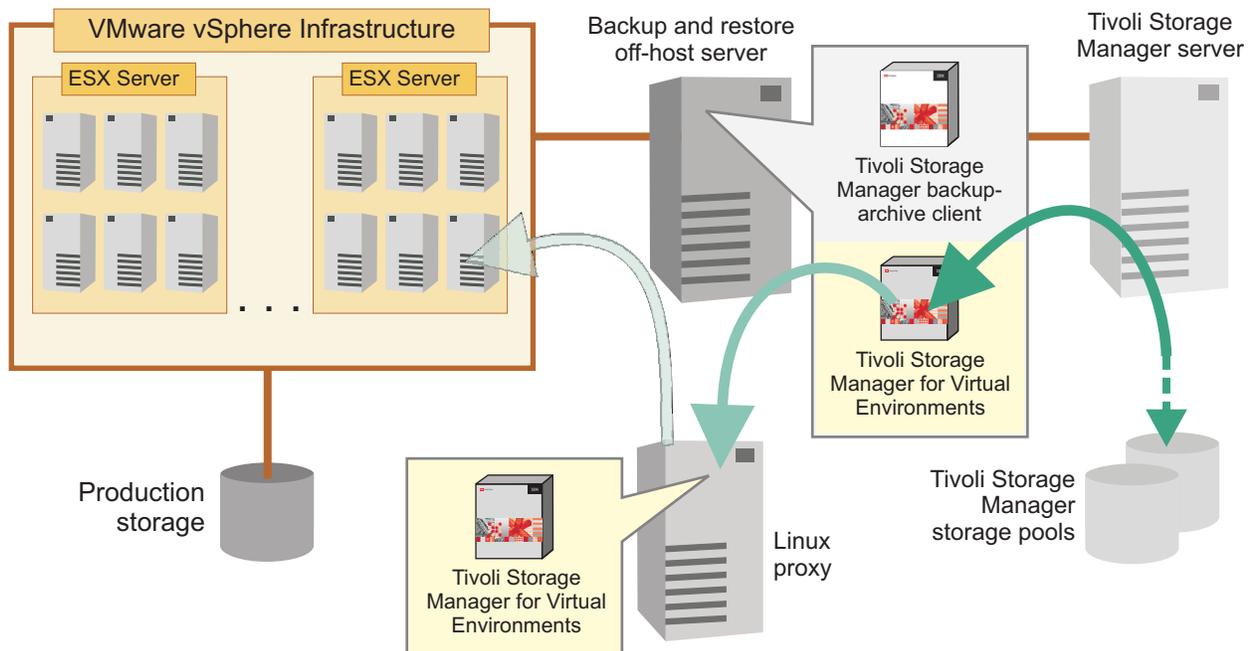


Figure 3. Off-host file-level restore for Linux

In-guest file-level restore and instant restore for Windows and for Linux

These configurations require Data Protection for VMware Recovery Agent to be installed in each virtual machine guest. The mount and instant restore processes are performed for a single partition from the backed up disk.

The Data Protection for VMware Recovery Agent node name would typically be granted access only to the virtual machine where it is running. To grant access, issue the following command from the backup-archive client node that owns the virtual machines:

```
set access backup "{\\VMFULL-vmdisplayname}\\*\\*" * mountnodename
```

The restore process is typically begun by a VMware user who logs in to the guest machine of the virtual machine.

For these configurations, be sure to compare the specific virtual machine guest operating system requirements with the supported levels of Data Protection for VMware Recovery Agent. If a specific operating system is not supported, determine if the off-host disk / block device exposure configuration can also be used for file-level recovery. Instant restore can only be used within a virtual machine guest for volumes that are not the operating system volume.

The data paths for in-guest file level restores are illustrated in Figure 4 on page 19 and Figure 5 on page 19. The data path for in-guest instant restore is illustrated in Figure 6 on page 20 and Figure 7 on page 20.

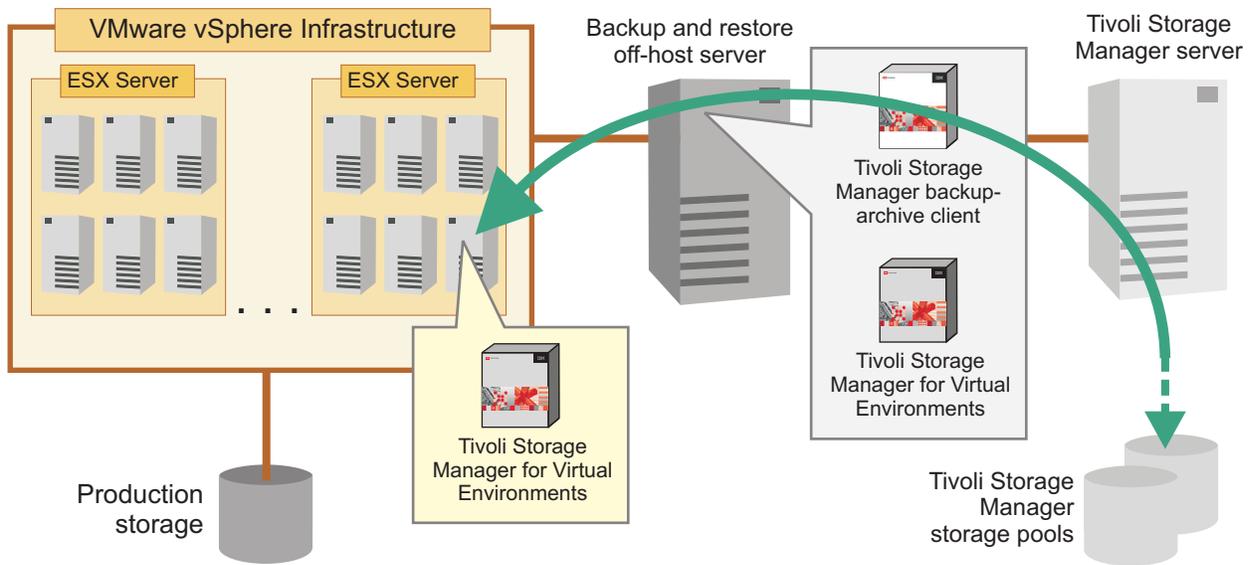


Figure 4. In-guest file-level restore for Windows

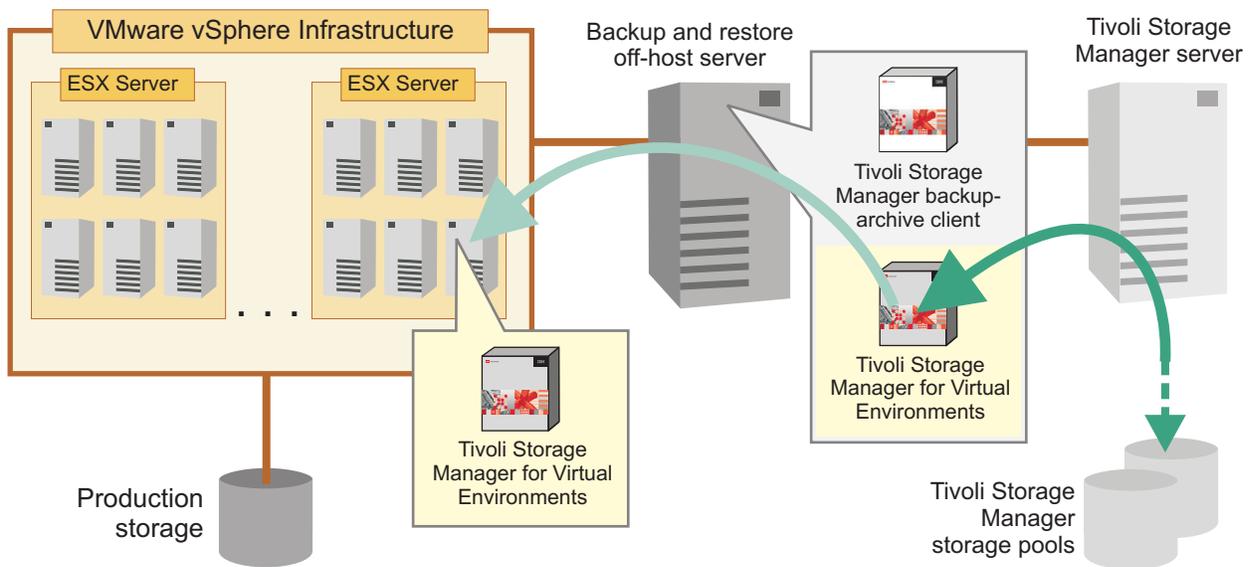


Figure 5. In-guest file-level restore for Linux

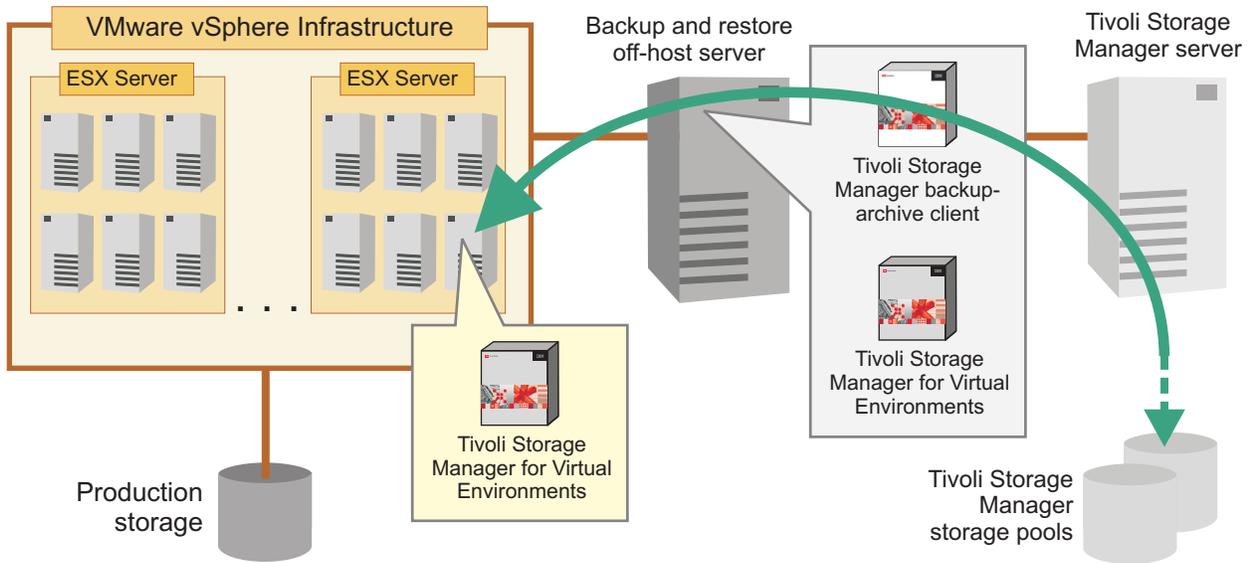


Figure 6. In-guest instant restore for Windows

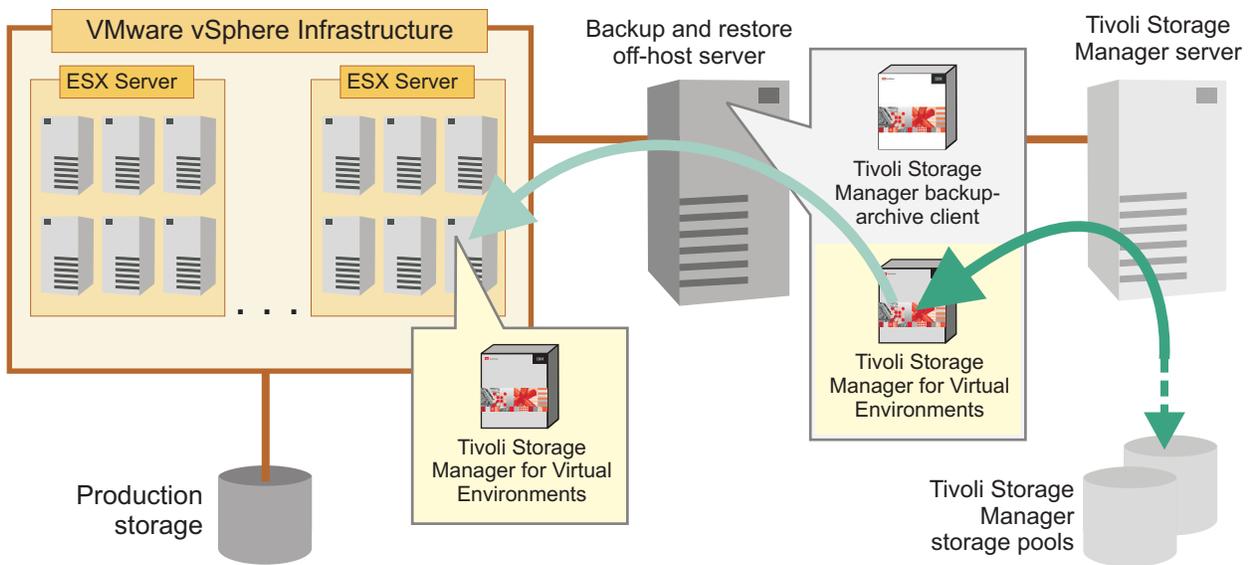


Figure 7. In-guest instant restore for Linux

Off-host iSCSI target

This configuration exposes an iSCSI target from the Windows instance of the off-host Data Protection for VMware Recovery Agent and manually uses an in-guest iSCSI initiator to access the disk snapshot. This configuration requires an iSCSI initiator to be installed within the virtual machine guest. This approach exposes an iSCSI LUN, rather than the off-host file-level restore for Windows and Linux, which expose an individual disk partition.

In this configuration, the mount process specifies the virtual machine guest iSCSI initiator name. After a disk snapshot has been mounted, it can be discovered and logged in to by using the iSCSI initiator in the virtual machine guest. This

centralized restore process would typically be initiated by a VMware administrator, Tivoli Storage Manager administrator or help desk personnel.

If you back up a virtual machine that contains GUID Partition Table (GPT) disks and want to mount the volume in the GPT disk, follow this procedure:

1. Mount the GPT disk as an iSCSI target.
2. Use the Microsoft iSCSI Initiator to log onto the target.
3. Open the Windows Disk Management to find the disk and bring it online. You can then view the volume in the GPT disk.

The data path for off-host iSCSI target device exposure is illustrated in Figure 8.

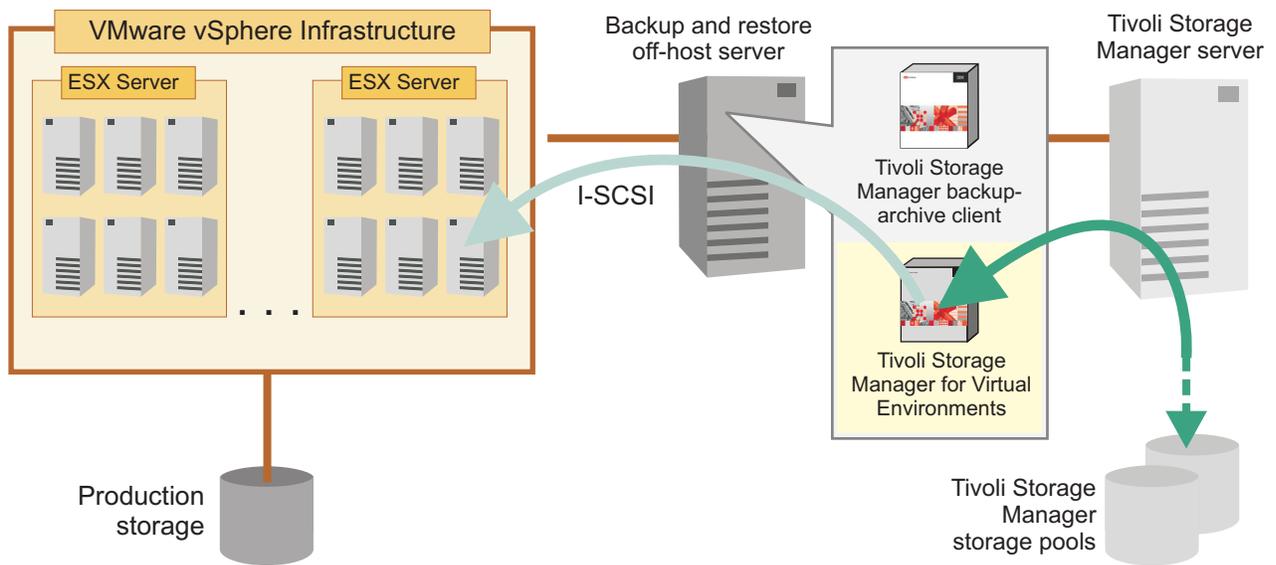


Figure 8. Off-host iSCSI target

Chapter 3. Installing, upgrading, and uninstalling Data Protection for VMware

Before installing or upgrading Data Protection for VMware, verify that your system meets all operating system, hardware, and software requirements.

For the system requirements, see the Chapter 2, “Planning,” on page 5 section. Each installation package presents you with an end user licensing file (EULA). If you do not accept the file, the installation stops.

Depending on your operating system environment, the following Data Protection for VMware components are available for installation:

Table 13. Available Data Protection for VMware components by operating system

Component	Linux	Windows
Data Protection for VMware Recovery Agent Provides virtual mount and instant restore capabilities.	√	√
Recovery Agent command-line interface Command-line interface used for mount operations (TDPVMwareShell.exe).		√
Documents Installs the quick start guide, readme file, and notices file.	√	√
Data Protection for VMware Enablement File Enables communication with Tivoli Storage Manager and used when creating incremental snapshots of your virtual machines.	√	√
Data Protection for VMware vCenter plug-in Integrates the product with the VMware vSphere client in order to back up, restore, and manage virtual machines in the vCenter. Includes the Data Protection for VMware command-line interface.	√ ¹	√ ²

Note:

1. Not available for 32-bit operating systems.
2. Not available for Microsoft Windows XP.

See “Supported operating systems” on page 5 for complete details regarding supported versions.

Preparing for installation

Before you begin installation, ensure that certain prerequisites are met.

- Remove any version of IBM Tivoli Storage Manager FastBack on the server. The Data Protection for VMware installation procedure checks for the existence of Tivoli Storage Manager FastBack on the server. If found, the installation fails.

- **Windows** Run the installation or upgrade process from a Windows logon ID with administrator authority.
- **Linux** Run the installation process as the root user. The root user profile must be sourced. If you use the **su** command to switch to root, use the **su -** command to source the root profile.
- **Linux** Ensure that the file `/etc/hosts` contains this text:
127.0.0.1 localhost
- When the installation path you plan to use contains non-English characters, you must install the following language packs before attempting to install any Data Protection for VMware component:
 - Appropriate operating system language pack
 - Appropriate Data Protection for VMware language pack
- If the embedded WebSphere Application Server (eWAS) is slow or busy, its folders and files are not removed during uninstallation. As a result, the existing eWAS blocks reinstallation of Data Protection for VMware. To prevent this issue, delete the eWAS folders (`C:\IBM\tivoli\tsm\tdpvmware\ewas`) before attempting to install Data Protection for VMware again.
- The Java™ Runtime Environment (JRE) is not part of the `setup.exe`, and it resides outside the install package under the `Extra` directory. In order to install the JRE, you must have the product DVD. If you did not install the JRE (it is installed if you install Mount), you need to run the `setup.exe` from the DVD so that the JRE can be copied to `C:\Program File\Tivoli\TSM\TDPVMware`.

Note: In non-English environments, the progress bars in the Data Protection for VMware installation screens contain English text. This issue is a known installation program limitation.

Installing Data Protection for VMware on Windows

Windows

You can install only one Data Protection for VMware vCenter plug-in on a machine. As a result, multiple Data Protection for VMware vCenter plug-ins are not allowed on the same machine.

To install Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. Start the installation program by running `setup.exe` for your operating system:
 - `CD\x86\setup.exe`
 - `CD\x64\setup.exe`
3. Choose the language to be used for the installation process and click **OK**.
4. The Welcome page opens. Click **Next**.
5. The Software License Agreement page opens. Read the terms of the license agreement. Select **I accept the terms in the license agreement** and click **Next**. If you do not accept the terms of the license agreement, the installation ends.
6. The Choose Destination Location page opens prompting you to specify where to install the software. You can accept the default location shown in the **Destination Folder** field, or click **Browse** to go to the location.
 - The default installation directory for the Data Protection for VMware Recovery Agent is `C:\Program Files\Tivoli\TSM\TDPVMware\mount`.

- The default installation directory for the Recovery Agent CLI is C:\Program Files\Tivoli\TSM\TDPVMware\shell.
- The installation directory for the Data Protection for VMware vCenter plug-in is C:\Program Files\Common Files\Tivoli\TDPVMware. This destination cannot be modified by the Choose Destination Location page.

Tip: When the Data Protection for VMware vCenter plug-in begins installing, clicking **Cancel** does not remove any files that were already installed. You must remove these files manually.

- The default installation directory for Windows 2008 and Windows Vista is C:\ProgramData\Tivoli\TSM\TDPVMware.
- The default installation directory for the embedded WebSphere Application Server (eWAS) is C:\IBM\tivoli\tsm\tdpvmware\ewas\.

Click **Next**.

- The Installation Type page opens, prompting you to select **Complete** or **Custom**:
 - Select **Complete** to install all of the components listed in Table 13 on page 23.
 - Select **Custom** to select only those components you want to install.

After completing your selection, click **Next** to continue.

- If you are installing the Data Protection for VMware Recovery Agent, the Define Data Protection for VMware Recovery Agent page opens. Enter the host name (or static IP address) where the Data Protection for VMware Recovery Agent is located. Click **Next**.
- A series of dialogues display that are used to register the Data Protection for VMware vCenter plug-in to the vCenter:

Tip: There might be a slight delay during registration when attempting to progress to the next dialog. This delay is due to Data Protection for VMware vCenter plug-in verifying and processing user input.

- Enter the vCenter Server IP address or name.
- Enter the vCenter user name.
- Enter the vCenter password. After entering the vCenter information, consider the following recommendations regarding the Derby Database and WebSphere Application Server configuration:
 - If you want to use default values for these applications, click **Next** for each dialog until the Installation Summary panel displays. Click **Install** to begin installing the files.

Note: On Windows x86, there is no Install button. Click **Next** in the last dialog to begin installing the files.

- If you want to modify the values for these applications, click **Advanced** and complete the following steps:
 - 1) Enter the TCP/IP port number for the Derby Database and click **Next**.
 - 2) Enter the default base port number for the WebSphere Application Server and click **Next**.

The Installation Summary panel displays. Click **Install** to begin installing the files. Click **Next** in the last dialog to begin installing the files.

Note: On Windows x86, the Installation Summary panel does not display.

Remember: Due to the installation and setup of multiple components, the installation process might take several minutes to complete.

If you restart the installation process after you have completed the initial installation, a Welcome window opens. From this window, there are three options:

- **Modify:** Use this option to select new program features or to delete installed features.
- **Repair:** Use this option to reinstall all program features installed during the previous setup.
- **Remove:** Use this option to uninstall Data Protection for VMware.

Installing Data Protection for VMware on Linux

Linux

You can install only one Data Protection for VMware vCenter plug-in on a machine. As a result, multiple Data Protection for VMware vCenter plug-ins are not allowed on the same machine.

Important: When the Deployment Engine, an IBM service component, is installed on your system, make sure that the hostname has not changed before installing Data Protection for VMware. If the host name has changed, the installation process might fail with this error: Deployment Engine failed to initialize. If the host name has changed, issue this command to correct the issue before installing Data Protection for VMware:

```
/usr/ibm/common/acs/bin/de_chghostname.sh
```

Restriction: Linux

- When installing the Data Protection for VMware vCenter plug-in on a system that contains a Tivoli Storage Manager API version earlier than 6.3.0, the installation fails and displays an error message. The earlier version of the Tivoli Storage Manager API does not support an upgrade. As a result, you must manually uninstall the existing Tivoli Storage Manager API. Then, upgrade the Tivoli Storage Manager API to version 6.3.0 (or later) before installing the Data Protection for VMware vCenter plug-in.
- When installing the Data Protection for VMware vCenter plug-in on a system that contains the Tivoli Storage Manager Backup-Archive Client, and its Tivoli Storage Manager API version is earlier than 6.3.0, the installation fails and displays an error message. In this situation, check the `/opt/tivoli/tsm/tdpvmware/TIVsm-API64RPM.log` file.

To install Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. Start the installation program by running `install-Linux.bin` for your platform. The default installation path for Linux is `/opt/tivoli/tsm/tdpvmware`.
3. Choose the language to be used for the installation process and click **OK**.
4. The Welcome page opens. Click **Next**.
5. The Software License Agreement page opens. Read the terms of the license agreement. Select **I accept the terms in the license agreement** and click **Next**. If you do not accept the terms of the license agreement, the installation ends.

6. The Choose Installation Folder page opens prompting you to specify where to install the software. You can accept the default location shown in the **Destination Folder** field, type the location name, or click **Browse** to go to the location. Click **Next**.
7. The Choose Install Set page opens, prompting you to select **Complete** or **Custom**:
 - Select **Complete** to install all of the components listed in Table 13 on page 23.
 - Select **Custom** to select only those components you want to install.
 After completing your selection, click **Next** to continue.
8. A series of dialogues display that configure the Data Protection for VMware vCenter plug-in to the vCenter:
 - If you want to use default values, select **Default** and click **Next** for each dialog until the Pre-Installation Summary panel opens. This panel contains a list of the settings you provided. Review the settings and click **Install** to begin installing the files.
 - If you want to modify the values, select **Customize** and click **Next**. Complete the following steps:
 - a. Enter a user name. A profile for this user name is created in `/home/<username>/tdpvmware/config`. Click **Next**.
 - b. Enter the Derby Database Port number and click **Next**.
 - c. Enter the WebSphere Application Server Default Base Port number and click **Next**.
 - d. Register the Data Protection for VMware vCenter plug-in by entering the following vCenter Server information:
 - Enter the vCenter Server IP address or name.
 - Enter the vCenter user name.
 - Enter the vCenter password.
 Click **Next**.
9. The Pre-Installation Summary panel opens. This panel contains a list of the settings you provided. Review the settings and click **Install** to begin installing the files.

Attention: If the following message is displayed at the end of installation, Data Protection for VMware installed successfully:

WARNING: A tool is either missing or is at an unsupported level. Check the logs in `opt/tivoli/tsm/tdpvmware/mount/engine/var`.

However, check the log file to determine which tool is missing or is not at the supported level and resolve the issue. The Data Protection for VMware Recovery Agent cannot function if the tool is missing or is not at the supported level.

Restriction: When there is not enough space to complete the installation, a warning message displays in the Preinstallation Summary page. This warning instructs you to exit the installation by clicking Cancel. However, Cancel is disabled. This issue is a known limitation. You can exit the installation by completing either of these tasks:

- In the Preinstallation Summary page, press **Alt + F4** to exit the installation immediately.
- Click **Previous**, then click **Cancel** (in the previous page) to exit the installation.

If you restart the installation process after you have completed the initial installation, a Choose Install Set window opens. From this window, there are two options:

- **Complete:** Use this option to install Data Protection for VMware.
- **Custom:** Use this option to customize the components to be installed.

Note:

If you were unable to install Data Protection for VMware successfully, see the "Manually removing Data Protection for VMware" procedure in "Uninstalling Data Protection for VMware a Linux system" on page 36.

Performing a clean installation of Data Protection for VMware on Linux

If a Linux installation is interrupted, you can usually restart it. However, if the installation fails to restart, a clean installation is required.

Before starting a clean installation, ensure that product is completely removed. Perform following steps to ensure a clean environment:

1. Remove all files from the failed installation:
 - a. Remove `<USER_INSTALL_DIR>`, which is the path where the failed installation was performed. For example: `/opt/tivoli/tsm/TDPVMware/`
 - b. Remove `<user.home>/IA-TDPVMware-00.log`. The `<username>` is the ID of the user who performed the installation.
2. Remove the Deployment Engine:
 - a. Remove the `/var/ibm/common` directory.
 - b. Remove the `/usr/ibm/common` directory.
 - c. Clean up the `/tmp` directory by removing the `acu_de.log` file, if it exists.
 - d. Remove the `/tmp` directory that contains the ID of the user that installed the Deployment Engine. ("root").
 - e. Remove all Deployment Engine entries from the `/etc/inittab` system file. The entries are delimited by `#Begin AC Solution Install` block and `#End AC Solution Install` block. Remove all text between those delimiters, and remove the delimiting text itself.
 - f. Remove all Deployment Engine references from the `/etc/services` system file.

After you have completed the previous steps, you can start the clean installation.

Installing language packs

Language packs can be installed after Data Protection for VMware is installed. The language packs are available on the product DVD.

Installing a language pack on Windows

You can install a Windows language pack after Data Protection for VMware has been installed. You can install one or more languages from the single package.

Data Protection for VMware supports installation of components on non-English versions of Windows, as well as non-ASCII objects (for example, host names, volume names, user names, passwords, and policies).

To install a language pack on a supported Windows operating system, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. From the Windows Start menu, select **Run** and enter the following command (where X represents the DVD drive letter or installation folder):
`X:\LanguagePacks\Windows\setup.exe`

Click **OK**

3. Follow the installation instructions contained in the prompt windows.
4. Click **Finish**.

Installing a language pack on Linux

You can install a Linux language pack after Data Protection for VMware is installed.

To install a language pack on a supported Linux operating system, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. Open a command prompt window and navigate to the /media directory. For example, type the `cd /media` command where /media represents the DVD mount point.
3. Run the installation process. For example, type the `./cdrom/TDPMware/LanguagePacks/Linux/installLP-Linux.bin` command. A Welcome page is displayed.
4. Follow the installation instructions contained in the windows.
5. Click **Finish**.

Installing Data Protection for VMware in silent mode

You can install Data Protection for VMware in the background. During this silent installation, no messages are displayed. After a silent installation completes, you must restart the system.

You can use the silent installation method for the following Data Protection for VMware components:

- Data Protection for VMware Recovery Agent
-  Command line
- Data Protection for VMware documents
-  Enablement file
- Data Protection for VMware vCenter plug-in

Use the procedure for your operating system:

- “Installing Data Protection for VMware on a Windows 32-bit system in silent mode” on page 30
- “Installing Data Protection for VMware on a Windows 64-bit system in silent mode” on page 30
- “Installing Data Protection for VMware on a Linux system in silent mode” on page 31

If you use the remote desktop to perform a silent install on a Windows XP system, the **Mount** icon does not appear in the task bar. To view and use the icon, you must manually restart the machine after installation is complete.

The Data Protection for VMware virtual volume kernel driver is not installed during the installation process. The virtual volume kernel driver is installed when Data Protection for VMware Recovery Agent is started for the first time.

Installing Data Protection for VMware on a Windows 32-bit system in silent mode

You can silently install Data Protection for VMware and the command line on a supported Windows 32-bit operating system:

Make sure that no embedded WebSphere Application Server (eWAS) folder exists before issuing a silent install operation.

To install Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. In the folder for Data Protection for VMware, go to the X86 folder.
3. In a text editor, open the `setup.iss` file.
4. Complete the following steps to edit the `setup.iss` file:
 - a. Locate the line that starts with the following string:
`szDir=`
 - b. (Optional) If you are not using the default installation path, edit this line to see the installation path that you are using.
 - c. Locate the line that starts with the following string:
`MOUNT_IP=`
 - d. Update the host name or IP address to refer to the Data Protection for VMware server that you have installed and are using.
 - e. Register the Data Protection for VMware vCenter plug-in to the vCenter by setting these three options:
`VCENTER_HOSTNAME=`: Enter the vCenter Server IP address or name.
`VCENTER_USERNAME=`: Enter the vCenter user name.
`VCENTER_PASSWORD=`: Enter the vCenter password.
 - f. Save and close the `setup.iss` file.
5. From a command prompt window, enter the following command:
`setup.exe /s /f1"<path_to_the_setup.iss_file>"`

Note: Specify an absolute path. Using a relative path can cause unpredictable results.

6. Restart the system.

Installing Data Protection for VMware on a Windows 64-bit system in silent mode

You can silently install Data Protection for VMware and the command line on a supported Windows 64-bit operating system.

Make sure that no embedded WebSphere Application Server (eWAS) folder exists before issuing a silent install operation.

To silently install Data Protection for VMware into the default location, follow these steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. In the folder for Data Protection for VMware, go to the X64 folder.
3. From a command prompt window, use the `cd` command to change the directory to the installation folder.
4. Enter the following command:

```
setup.exe /s /v"/qn REBOOT=ReallySupress
VCENTER_HOSTNAME=<vCenter hostname or IP>
VCENTER_USERNAME=<vCenter user name>
VCENTER_PASSWORD=<vCenter password>"
```
5. Restart the system.

Performing a clean installation in a non-default location

To perform a clean installation in a non-default location, follow these steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. In the folder for Data Protection for VMware, go to the X64 folder.
3. From a command prompt window, use the `cd` command to change directory to the installation folder.
4. Enter the following command:

```
setup.exe/s /v"/qn REBOOT=ReallySupress
INSTALLDIR="\<your_installation_directory>\\"
VCENTER_HOSTNAME=<vCenter hostname or IP>
VCENTER_USERNAME=<vCenter user name>
VCENTER_PASSWORD=<vCenter password>"
```

Note: This command must be entered on one line. This example shows five lines to accommodate page formatting.
5. Restart the system.

Installing Data Protection for VMware on a Linux system in silent mode

You can silently install Data Protection for VMware on a supported Linux operating system. Installation on Linux does not include the command line or the enablement file.

By default, Data Protection for VMware installs the Data Protection for VMware vCenter plug-in. However, in order to successfully install the Data Protection for VMware vCenter plug-in, you must first specify values for these three options in the `installer.properties` file:

- `VCENTER_HOSTNAME=`: Enter the vCenter Server IP address or name.
- `VCENTER_USERNAME=`: Enter the vCenter user name.
- `VCENTER_PASSWORD=`: Enter the vCenter password.

If you do not want to install the Data Protection for VMware vCenter plug-in, edit the `installer.properties` file and remove these three options.

Make sure that no embedded WebSphere Application Server (eWAS) folder exists before issuing a silent install operation.

Restriction: Linux

- When installing the Data Protection for VMware vCenter plug-in on a system that contains a Tivoli Storage Manager API version earlier than 6.3.0, the installation fails and displays an error message. The earlier version of the Tivoli Storage Manager API does not support an upgrade. As a result, you must manually uninstall the existing Tivoli Storage Manager API. Then, upgrade the Tivoli Storage Manager API to version 6.3.0 (or later) before installing the Data Protection for VMware vCenter plug-in.
- When installing the Data Protection for VMware vCenter plug-in on a system that contains the Tivoli Storage Manager Backup-Archive Client, and its Tivoli Storage Manager API version is earlier than 6.3.0, the installation fails and displays an error message. In this situation, check the `/opt/tivoli/tsm/tdpvmware/TIVsm-API64RPM.log` file.

To install Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. Uncomment this entry in the `installer.properties` file in order to accept the license:

```
#LICENSE_ACCEPTED=TRUE
```

As a result, the entry must be exactly as shown in the following example:

```
LICENSE_ACCEPTED=TRUE
```

3. Open the Linux folder, which is located in the Data Protection for VMware folder and choose one of the following installations:
 - a. For the default installation, enter the following command in the command prompt window:

```
./install-Linux.bin -i silent -DLICENSE_ACCEPTED=true  
-DVCENTER_HOSTNAME=<vCenter hostname or IP>  
-DVCENTER_USERNAME=<vCenter user name>  
-DVCENTER_PASSWORD=<vCenter password>
```
 - b. For a custom installation, enter the following command into the command prompt window:

```
./install-Linux.bin -i silent -f <full_path> <properties_file_name>
```

Upgrading Data Protection for VMware

Information about how to upgrade Data Protection for VMware is available.

You must have administrator privileges to upgrade your Windows or Linux Data Protection for VMware. An upgrade cannot install new components. For example, if you have installed only the Data Protection for VMware Recovery Agent component, an upgrade cannot install the command-line component. In such a scenario, you must run the installation program again and then select the missing component to install.

The Data Protection for VMware Recovery Agent on Linux version of the Data Protection for VMware Recovery Agent must be the version as the Windows Data Protection for VMware Recovery Agent that is used as the proxy. Therefore, when you upgrade the Data Protection for VMware Recovery Agent on Linux version of

the Data Protection for VMware Recovery Agent, you must also upgrade the Windows version of the Data Protection for VMware Recovery Agent.

Before proceeding with the upgrade, make sure there are no active instant restore or mount sessions. Also, make sure the existing Data Protection for VMware Recovery Agent is closed.

Important: When upgrading from Data Protection for VMware 6.2 to 6.3, in order to install the Data Protection for VMware vCenter plug-in, you must run the setup.exe (Windows) or install-Linux.bin (Linux) file again, after the initial upgrade procedure completes. During the subsequent upgrade operation:

- **Windows** Select Modify, then select VMware vCenter plug-in.
- **Linux** Go to the Choose Product Features panel, select Customize, then select VMware vCenter plug-in.

You can install only one Data Protection for VMware vCenter plug-in on a machine. As a result, multiple Data Protection for VMware vCenter plug-ins are not allowed on the same machine. The Data Protection for VMware vCenter plug-in is not available on some versions of Windows and Linux. See “Supported operating systems” on page 5 for supported versions.

To upgrade Data Protection for VMware, complete the following steps:

1. Download the code package.
2. From the folder where you saved the code package start the upgrade process:
 - a. **Windows:** Run the setup.exe file.
 - b. **Linux:** Run the install-Linux.bin file.
3. **Windows:** A pop-up message displays this text: "The Existing Data Protection for VMware is going to be upgraded."
4. If you confirm the upgrade, the installer updates the files.
5. **Linux:** After completing the upgrade, if the /usr/ibm/common/acsi/bin/listIU.sh command fails to show the installed units, complete these steps:
 - a. Edit the setenv.sh file. For example:

```
#vi /var/ibm/common/acsi/setenv.sh
```
 - b. Add the Deployment Engine JRE path to the SI_EXT_DIRS variable. For example:

```
SI_EXT_DIRS=/usr/ibm/common/acsi/jre/lib/ext:/usr/ibm/common/acsi/lib
```
 - c. Save the setenv.sh file.

Upgrading Data Protection for VMware on a Windows 32-bit system in silent mode

You can silently upgrade Data Protection for VMware on a supported 32-bit operating system

To upgrade Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware DVD into the DVD drive.
2. In the Data Protection for VMware folder, go to the X86 folder.
3. In a text editor, open the upgrade.iss file.
4. Edit the upgrade.iss file:
 - a. Locate the line that starts with the following string: szDir=

- b. **Optional:** If you are not using the default installation path, edit this line to refer to the installation path that you are using.
- c. Save and close the upgrade.iss file.
5. From a command prompt window, enter the following command:

```
setup.exe /s /f1"<path_to_the_upgrade.iss_file>"
```
6. Restart the system.

Upgrading Data Protection for VMware on a Windows 64-bit system in silent mode

You can silently upgrade Data Protection for VMware on a supported 64-bit operating system.

To upgrade Data Protection for VMware, complete the following steps:

1. Either download the code package or insert the Data Protection for VMware product DVD into the DVD drive.
2. In the folder for Data Protection for VMware, either go to the X64 folder.
3. From a command prompt window, enter the following command:

```
setup.exe /s /v"/qn REBOOT=ReallySuppress  
INSTALLDIR="\<path_to_the_install_directory>"
```

Note: This command must be entered on one line. This example shows two lines to accommodate page formatting.

Upgrading Data Protection for VMware on a Linux system in silent mode

You can silently upgrade Data Protection for VMware on a supported Linux operating system.

Restriction: Linux

- When installing the Data Protection for VMware vCenter plug-in on a system that contains a Tivoli Storage Manager API version earlier than 6.3.0, the installation fails and displays an error message. The earlier version of the Tivoli Storage Manager API does not support an upgrade. As a result, you must manually uninstall the existing Tivoli Storage Manager API. Then, upgrade the Tivoli Storage Manager API to version 6.3.0 (or later) before installing the Data Protection for VMware vCenter plug-in.
- When installing the Data Protection for VMware vCenter plug-in on a system that contains the Tivoli Storage Manager Backup-Archive Client, and its Tivoli Storage Manager API version is earlier than 6.3.0, the installation fails and displays an error message. In this situation, check the `/opt/tivoli/tsm/tdpvmware/TIVsm-API64RPM.log` file.

To upgrade Data Protection for VMware, complete the following steps:

1. Either download the code package, or insert the Data Protection for VMware product DVD into the DVD drive.
2. From the folder for Data Protection for VMware go to the Linux folder.
3. From a command prompt window, enter the following command:

```
./install-Linux.bin -i silent -DLICENSE_ACCEPTED=true
```

Uninstalling Data Protection for VMware

The process for uninstalling Data Protection for VMware is the same for a new installation and for an upgraded version.

Restriction: You must unmount all virtual volumes before uninstalling Data Protection for VMware Recovery Agent. Otherwise, these mounted virtual volumes cannot be unmounted after Data Protection for VMware Recovery Agent is reinstalled.

If the embedded WebSphere Application Server (eWAS) is slow or busy, its folders and files are not removed during uninstallation. As a result, the existing eWAS blocks reinstallation of Data Protection for VMware. To prevent this issue, delete the eWAS folders (C:\IBM\tivoli\tsm\tdpvmware\ewas) before attempting to install Data Protection for VMware again.

1. Start the uninstall process:
 - **Windows:** Either run the setup.exe file, or select **Add or Remove Programs** from the Windows Control Panel.
 - **Linux:** Run the install-Linux.bin file.
2. A screen opens. Select **Modify, Repair, or Remove**
3. Select **Remove**. Data Protection for VMware is completely uninstalled.
4. For a Windows uninstallation of Data Protection for VMware Recovery Agent, you are asked to reboot the computer, in order to complete the uninstallation of Data Protection for VMware Recovery Agent drivers. You can choose to reboot later.

Uninstalling Data Protection for VMware for a Windows 32-bit system in silent mode

You can silently uninstall Data Protection for VMware on a supported Windows 32-bit operating system.

To uninstall Data Protection for VMware, complete the following steps:

1. In the installation directory for Data Protection for VMware, go to the X86 folder.
 2. Identify the Data Protection for VMware vCenter plug-in by setting these three options in the uninstall.iss file:
 - VCENTER_HOSTNAME=: Enter the vCenter Server IP address or name.
 - VCENTER_USERNAME=: Enter the vCenter user name.
 - VCENTER_PASSWORD=: Enter the vCenter password.
 3. From a command prompt window, enter the following command:

```
setup.exe /s /f1"  
<full absolute path_to_the_uninstall.iss_file and uninstall.iss>
```
- Note:** The command must be entered on one line. This example shows two lines to accommodate page formatting.
4. Restart the system.

Uninstalling Data Protection for VMware for Windows 64-bit system in silent mode

You can silently uninstall Data Protection for VMware on a supported Windows 64-bit operating system.

To uninstall Data Protection for VMware, complete the following steps:

1. In the installation directory for Data Protection for VMware, go to the X64 folder.
2. From a command prompt window, enter the following command to uninstall Data Protection for VMware:

Note: This command also unregisters the Data Protection for VMware vCenter plug-in.

```
setup.exe /s /v"/qn REBOOT=ReallySuppress REMOVE=ALL  
VCENTER_HOSTNAME=<vCenter hostname or IP>  
VCENTER_USERNAME=<vCenter user name>  
VCENTER_PASSWORD=<vCenter password>"
```

To uninstall Data Protection for VMware and skip Data Protection for VMware vCenter plug-in unregistration, enter this command:

```
setup.exe /s /v"/qn REBOOT=ReallySuppress REMOVE=ALL  
IGNORE_VCENTER_UNREGISTER=1"
```

3. Restart the system.

Uninstalling Data Protection for VMware a Linux system

You can uninstall Data Protection for VMware on a supported Linux operating system.

When you uninstall Data Protection for VMware on a Linux system, by default, the type of uninstallation is the same process as the type of original installation. To use a different uninstallation process, specify the correct parameter. For example, if you used a silent installation process, you can use the installation wizard to uninstall by specifying the `-I` swing parameter. Run the uninstallation process as the root user. The root user profile must be sourced. If you use the `su` command to switch to root, use the `su -` command to source the root profile.

When the uninstall process begins removing program files, canceling the uninstall process does not return the system to a clean state. This situation might cause the reinstallation attempt to fail. As a result, clean the system by completing the tasks described in "Manually removing Data Protection for VMware" on page 37.

To uninstall Data Protection for VMware, complete the following steps:

1. Change to the directory for the uninstallation program. The following path is the default location to the uninstallation program: `/opt/tivoli/tsm/TDPVMware/_uninst/TDPVMware/`
2. Depending on the type of installation, use one of the following methods to uninstall Data Protection for VMware:

Note: The commands in this procedure must be entered on one line. These examples show two lines to accommodate page formatting.

- To use the installation wizard to uninstall Data Protection for VMware, enter this command:

```
./ Uninstall_IBM\ Tivoli\ Storage\ Manager\ for\ Virtual\ Environments\ Data\  
Protection\ for\ VMware -I swing
```

- To use the console to uninstall Data Protection for VMware, enter this command:

```
./ Uninstall_IBM\ Tivoli\ Storage\ Manager\ for\ Virtual\ Environments\ Data\  
Protection\ for\ VMware -i console
```

- To silently uninstall Data Protection for VMware, enter this command:

```
./ Uninstall_IBM\ Tivoli\ Storage\ Manager\ for\ Virtual\ Environments\ Data\ Protection\ for\ VMware -i silent -f uninstall.properties
```

The `uninstall.properties` file contains the vCenter connection information. This information is needed to uninstall the Data Protection for VMware vCenter plug-in.

Manually removing Data Protection for VMware

When Data Protection for VMware cannot be installed using the standard installation procedure, you must manually remove Data Protection for VMware from the system as described in these steps. Complete this process as the root user.

1. If you installed the Data Protection for VMware vCenter plug-in, remove its package from the Package Manager database with this command:

```
rpm -e TIVsm-TDPVMwarePlugin
```
2. Remove the product entries from the Deployment Engine:
 - a. Issue this command to view a list of all entries:

```
/usr/ibm/common/acs/bin/de_lsrootiu.sh
```
 - b. Issue this command to remove the installed unit entries that are related to Data Protection for VMware:

```
/usr/ibm/common/acs/bin/deleteRootIU.sh <UUID> <discriminant>
```
3. Back up the global registry file (`/var/.com.zerog.registry.xml`). After the file is backed up, remove all tags related to Data Protection for VMware.
4. Remove all files located in the installation directory (`/opt/tivoli/tsm/tdpvmware`). Also remove any shortcuts located on the desktop.
5. Back up the log files located in the root directory that contain TDPVMware in the file name. For example, `IA-TDPVMware-00.log` or `IA-TDPVMware_Uninstall-00.log`. Remove these log files after they are backed up. By removing them, you can view any error issued if the installation process fails again.
6. Attempt to install the product again as described in “Installing Data Protection for VMware on Linux” on page 26.

Chapter 4. Starting and configuring Data Protection for VMware

This section provides instructions for starting and configuring Data Protection for VMware.

Starting and running services for Data Protection for VMware

By default, when you start the Windows operating system, Data Protection for VMware Recovery Agent is started under the Local System Account.

Data Protection for VMware Recovery Agent services on Windows 7, Windows Vista, and Windows 2008

When you start the Data Protection for VMware Recovery Agent from the Windows Start menu, the service is automatically stopped. When the Data Protection for VMware Recovery Agent, started from the Start menu finishes, the service starts automatically. In addition, for these operating systems, the service does not provide a GUI. In order to use the GUI, go to the Windows Start menu and select **All Programs > Tivoli Storage Manager > Data Protection for VMware > Data Protection for VMware Recovery Agent**.

WebSphere Application Server (eWAS) service

You can verify that the IBM WebSphere Application Server (eWAS) service is running by completing the following task:

- **Windows** Go to **Start > Control Panel > Administrative Tools > Services** and verify that the status of IBM WebSphere Application Server V7.0 - TSMVPlugin is Started.
- **Linux** Go to the scripts directory (`/opt/tivoli/tsm/tdpvmware/common/scripts/`) and issue this command:

```
./ewas status
```

These init scripts can also be used to stop and start the eWAS service:

```
/etc/init.d/ewas stop  
/etc/init.d/ewas start
```

Tivoli Storage Manager client acceptor

You can verify that the Tivoli Storage Manager client acceptor is running by completing the following task:

Windows Go to **Start > Control Panel > Administrative Tools > Services** and verify that the status of Data Protection for VMware command-line interface is Started.

Linux Go to the scripts directory (`/opt/tivoli/tsm/tdpvmware/common/scripts/`) and issue this command:

```
./vmclid status
```

- If the daemon is not running, issue this command to manually start the daemon:

```
/opt/tivoli/tsm/tdpvmware/common/scripts/vmcli --daemon
```

These init scripts can also be used to stop and start the daemon:

```
./vmclid stop  
./vmclid start
```

Configuration and log files

Configuration files are required for Data Protection for VMware to run correctly.

Windows In the Windows operating system, logs are placed as follows:

Windows 2003 and XP: C:\Documents and Settings\All Users\Application Data\Tivoli\TSM\tdpvmware\
Data\Tivoli\TSM\tdpvmware\
Data\Tivoli\TSM\tdpvmware\

Windows 7, Vista, and 2008: C:\ProgramData\Tivoli\TSM\TDPVMware

In these paths there are subdirectories with a folder for each Data Protection for VMware service. For example, the subdirectories contain folders labeled mount and shell.

Linux In the Linux operating system logs are placed in <user.home>/tivoli/tsm/ve/mount/log. In addition, logs are placed in /opt/tivoli/tsm/TDPVMware/mount/engine/var

The log file with the most recent data is stored in the log file with the *040* number. When a log file reaches the maximum size limit, a new log file is created. The log file name is the same except that the log file number decrements by one. Specifically, the data in the log file with the *040* number is copied to a log file with the *039* number. The log file with the *040* number contains the newest log file data. When *040* again reaches maximum file size, the *039* file contents move to *038* and the *040* information goes to *039* again.

Note: Although log files have extensions of `"*.sf"`, they are plain text files.

Chapter 5. Using the Data Protection for VMware vCenter plug-in

The Data Protection for VMware vCenter plug-in is the primary interface for controlling backup, restore and reporting operations in your VMware environment.

Complete all tasks described in “Configuring the Data Protection for VMware vCenter plug-in” before attempting any operations.

Start the Data Protection for VMware vCenter plug-in by clicking the icon in the Solutions and Applications window of your vSphere Client.

The Data Protection for VMware vCenter plug-in is installed on a VM guest machine, off-host physical machine, or the same machine used by the Tivoli Storage Manager Administration Center. Each Data Protection for VMware vCenter plug-in installation manages a backup domain that contains one or more VMware data centers. By default, a domain contains all data centers that are associated with a vCenter. You can use the Data Protection for VMware vCenter plug-in to limit a domain to one or more VMware data centers.

Configuring the Data Protection for VMware vCenter plug-in

You must configure and verify each component so that the Data Protection for VMware vCenter plug-in is ready for backup and restore operations.

The Data Protection for VMware vCenter plug-in and Data Protection for VMware command-line interface communicate with Tivoli Storage Manager by using proxy node relationships. A Tivoli Storage Manager server must be available to register the nodes.

Complete each of these tasks before attempting any Data Protection for VMware vCenter plug-in operation. In the following tasks, example node names are used to illustrate the relationships among the various nodes.

1. “Register Tivoli Storage Manager nodes”
2. “Setting up the Tivoli Storage Manager Backup-Archive Client on the vStorage Backup Server for each datamover node” on page 44
3. “Configuring the Data Protection for VMware command-line interface” on page 46
4. “Start the Data Protection for VMware vCenter plug-in” on page 48

Register Tivoli Storage Manager nodes

The Data Protection for VMware vCenter plug-in uses Tivoli Storage Manager nodes and their proxy relationships to complete backup and restore operations.

Review this information in order to understand the Tivoli Storage Manager nodes that are used by the Data Protection for VMware vCenter plug-in:

Table 14. Tivoli Storage Manager nodes

Node	Description	Example in procedure
vCenter node	The virtual node that represents a vCenter.	VC1

Table 14. Tivoli Storage Manager nodes (continued)

Node	Description	Example in procedure
data center node ¹	The virtual node that maps to a data center. The data center nodes hold the data.	VC1_DC1 VC1_DC2
Data Protection for VMware command-line interface node	The node that connects the Data Protection for VMware command-line interface to the Tivoli Storage Manager server and the Tivoli Storage Manager data mover node.	VC1_VCLI1
Tivoli Storage Manager data mover node	The node name for the Tivoli Storage Manager Backup-Archive Client that is installed on the vStorage Backup Server. This node performs the data movement. You can have multiple Tivoli Storage Manager data mover nodes for each vStorage Backup Server. This feature is useful when multiple schedules run in parallel on a single server. In this scenario, each node corresponds to a scheduler instance.	VC1_DC1_DM1 VC1_DC2_DM2

1. Certain restrictions apply to the data center node when the data center name is changed. See the description of VE_DATACENTER_NAME in “Profile parameters” on page 64 for complete information. In addition, the Data Protection for VMware vCenter plug-in does not support data centers with the same name in the vCenter.

In the following illustration, the arrow points from the proxy agent node to the proxy target node.

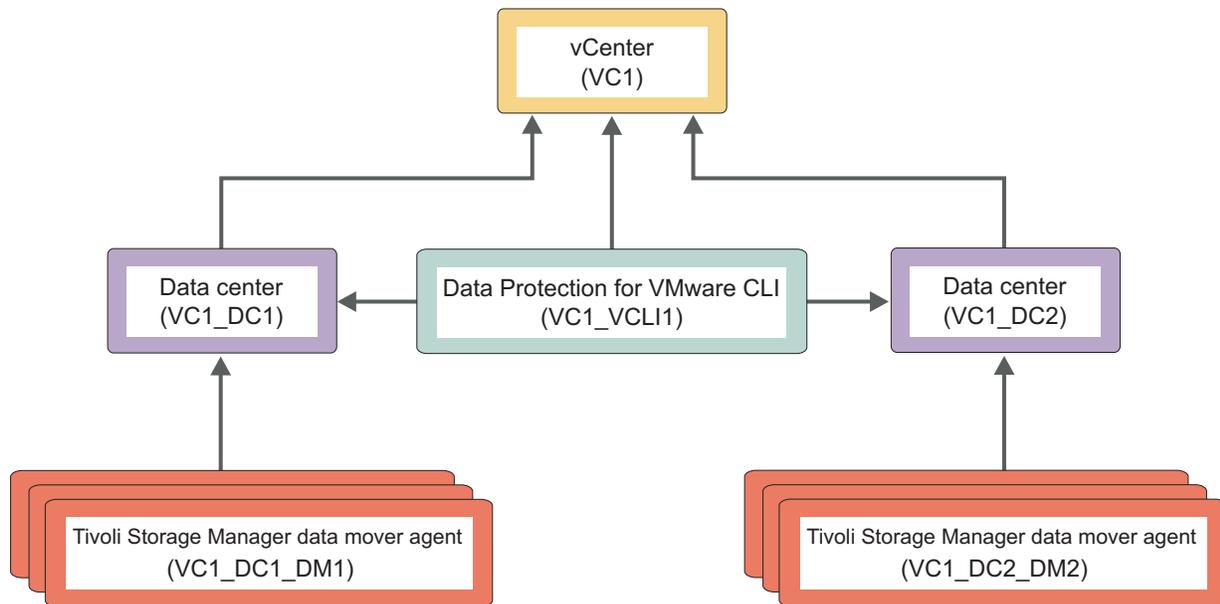


Figure 9. Proxy relationships among the nodes

In this illustration, the nodes are presented in context of an actual production environment:

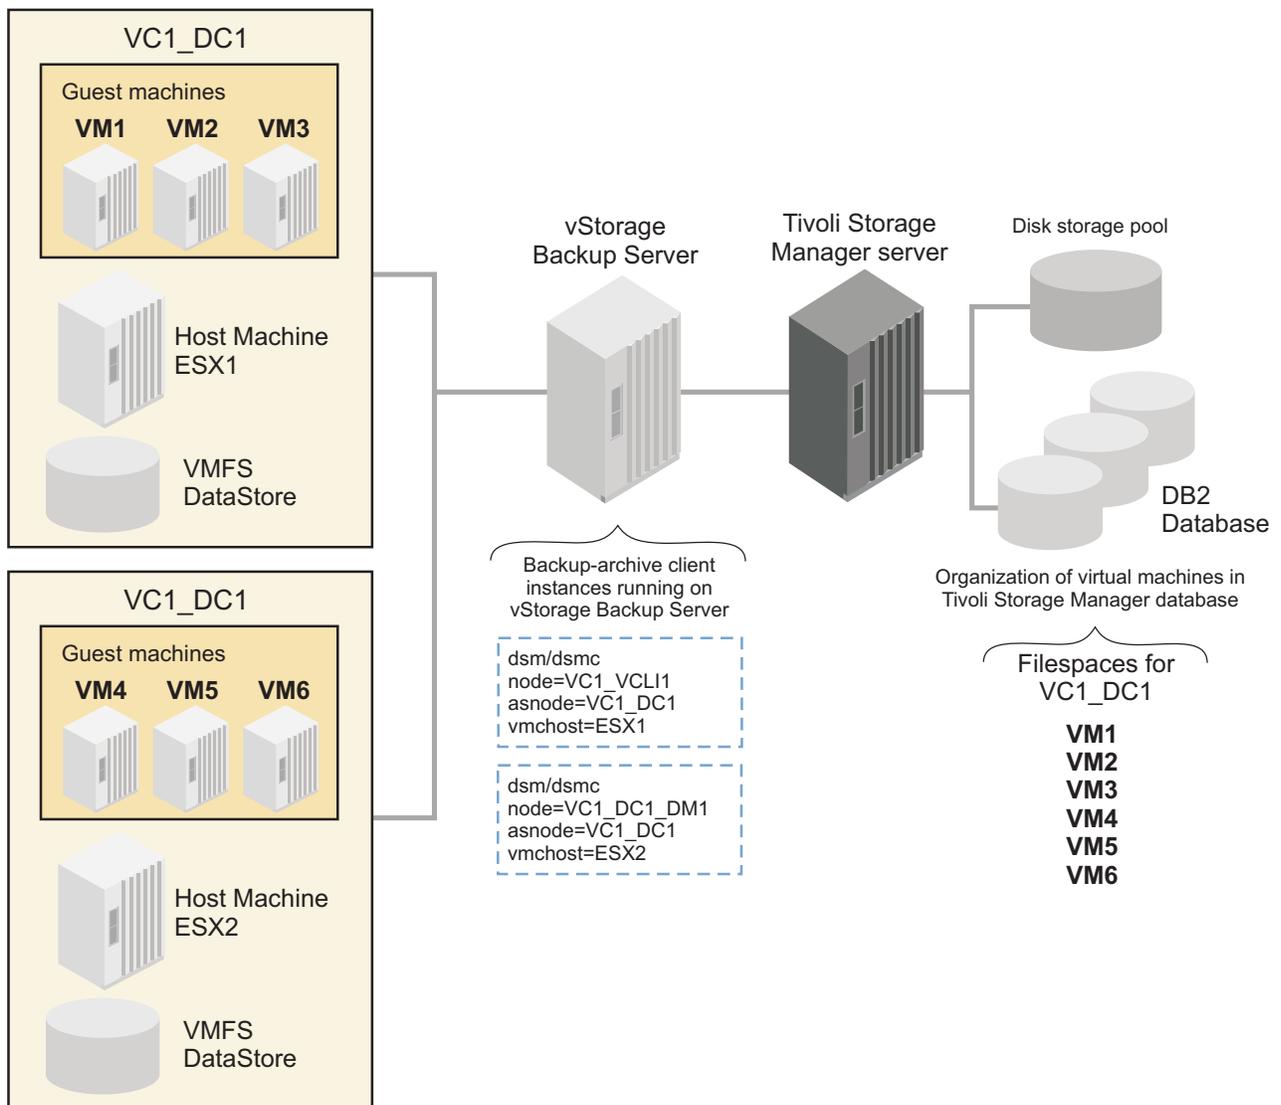


Figure 10. Tivoli Storage Manager nodes

Tip: All steps in this procedure are completed on the Tivoli Storage Manager server.

1. Register the following nodes:

- VC1
- VC1_DC1
- VC1_DC2
- VC1_VCLI1
- VC1_DC1_DM1
- VC1_DC2_DM2

2. Define proxy relationships for these nodes:

a. Grant proxy authority to the vCenter node by issuing this command:

```
grant proxynode target=VC1 agent=VC1_DC1,VC1_DC2,VC1_VCLI1
```

This command grants the VC1_DC1, VC1_DC2, and VC1_VCLI1 nodes the authority to perform tasks on behalf of the VC1 node.

- b. Grant proxy authority to the data center node by issuing this command:

```
grant proxynode target=VC1_DC1 agent=VC1_VCLI1,VC1_DC1_DM1
```

This command grants the VC1_VCLI1 and VC1_DC1_DM1 nodes the authority to perform tasks on behalf of the VC1_DC1 node.

- c. Grant proxy authority to the other data center node by issuing this command:

```
grant proxynode target=VC1_DC2 agent=VC1_VCLI1,VC1_DC2_DM2
```

This command grants the VC1_VCLI1 and VC1_DC2_DM2 nodes the authority to perform tasks on behalf of the VC1_DC2 node.

- Tivoli Storage Manager data mover node VC1_DC1_DM1 now has authority to move data for data center node VC1_DC1.
- Tivoli Storage Manager data mover node VC1_DC2_DM2 now has authority to move data for data center node VC1_DC2.

Tip: Grant proxy authority to all of the data center nodes in your vCenter.

Important: Whenever a node is renamed, the administrator node (associated with this renamed node) must also be renamed to match the renamed node. If the renamed node does not have an administrator node with the same name, remote access fails. See *The inquire_detail command failed with Return Code 53* in the “Troubleshooting” on page 54 section for instructions regarding how to complete this task.

Setting up the Tivoli Storage Manager Backup-Archive Client on the vStorage Backup Server for each datamover node

Update the Tivoli Storage Manager client options file (dsm.opt) and verify connectivity on the vStorage Backup Server.

This procedure describes how to set up the backup-archive client on the vStorage Backup Server and verify that you can back up a virtual machine.

Tip: All steps in this procedure are completed on the vStorage Backup Server.

1. Update the backup-archive client dsm.opt options file with these settings:

- **Windows** Specify these options in the dsm.opt options file.
- **Linux** Specify these options in the dsm.sys file, in the stanza for the Tivoli Storage Manager data mover node.

PASSWORDACCESS

Specify GENERATE so that the password is generated automatically (instead of a user prompt).

VMCHOST

Specify the host name of the vCenter (or ESX server) where the off-host backup commands are directed.

VMCUSER

Specify the user name of the vCenter (or ESX server) where the off-host backup commands are directed.

VMFULLTYPE

Specify VSTOR. This setting designates that the VMware vStorage API for Data Protection is used to run the backup.

VMBACKUPTYPE

Specify FULLVM. This setting designates that a full VM backup is run. This value is necessary to run full VM and full VM incremental backups.

VMCPW

Use the **dsmc set password** command or backup-archive client GUI Preferences Editor to set the password of the vCenter (or ESX server) user name. For example:

```
dsmc set password -type=vm vcenter vmcadmin vmadminpwd
```

This command saves your vCenter credentials. As a result, any node can use the credentials on the vStorage Backup Server by setting the VMCHOST and VMCUSER options. By using this method, the password is not shown in the options file (as shown in the dsm.opt file example).

An example dsm.opt file with these settings is provided here:

```
PASSWORDAccess generate
VMCHost vcenter.storage.usca.example.com
VMCUser vmcadmin
VMFULLTYPE VSTOR
VMBACKUPTYPE FULLvm
```

2. Specify the **-asnodename** parameter when starting the backup-archive client command-line interface:

```
dsmc -asnodename=VC1_DC1
```

Make sure that after your initial sign-on, you are not prompted for your password.
Complete these tasks:
 - a. Verify the connection to the Tivoli Storage Manager server by issuing this command:

```
dsmc query session
```

This command shows information about your session, including the current node name, when the session was established, server information, and server connection information.
 - b. Verify you can back up a virtual machine by issuing this command:

```
dsmc backup vm vm1
```

In Steps 2b and 2d, vm1 is the name of the virtual machine.
 - c. Verify that the backup completed successfully by issuing this command:

```
dsmc query vm "*"
```
 - d. Verify that the virtual machine can be restored by issuing this command:

```
dsmc restore vm vm1 -vmname=vm1-restore
```
3. Set up the client acceptor and scheduler by completing these tasks:
 - **Windows** Start the Tivoli Storage Manager Client Configuration wizard on the system where the Data Protection for VMware vCenter plug-in is installed:
 - a. Select Help me configure the TSM Web Client. Enter the information as prompted. When completed successfully, return to the wizard welcome page and proceed to Step b.
 - b. Select Help me configure the TSM Client Scheduler. Enter the information as prompted.
 - **Linux** Specify these options in the dsm.sys file, in the stanza for the Tivoli Storage Manager data mover node:

- a. Specify the `managementservices` option with these two parameters:
`managementservices schedule webclient`

This setting directs the client acceptor to manage both the Web client and the scheduler.

- b. (Optional) If you want to direct schedule and error information to log files other than the default files, specify the `schedlogname` and `errorlogname` options with the fully qualified path and file name in which to store log information. For example:

```
schedlogname /vmsched/dsmsched_dm.log
errorlogname /vmsched/dsmerror_dm.log
```

4. Verify that the client acceptor and agent are set up correctly:
 - a. Log on to a remote system.
 - b. Use a web browser to connect to the HOST1 system by using this address and port:
`http://HOST1.xyz.yourcompany.com:1581`

Tip: When the IP address changes on the system where the Data Protection for VMware vCenter plug-in is installed, you must set up the client acceptor again (Step 3) so that the Data Protection for VMware vCenter plug-in becomes enabled for operations. Otherwise, the Plug-in Manager shows the Data Protection for VMware vCenter plug-in status as disabled.

Configuring the Data Protection for VMware command-line interface

Update the Data Protection for VMware command-line interface profile on the system where the Data Protection for VMware vCenter plug-in is installed.

The profile (`vmcliprofile`) is located in this directory on the system where the Data Protection for VMware vCenter plug-in is installed:

```
Linux /home/username/tdpvmware/config
```

```
Windows 32-bit: C:\Program Files\Common Files\Tivoli\TDPVMware\VMwarePlugin\scripts
```

```
Windows 64-bit: C:\Program Files (x86)\Common Files\Tivoli\TDPVMware\VMwarePlugin\scripts
```

Tip: All steps in this procedure are completed on the system where the Data Protection for VMware vCenter plug-in is installed.

1. Update the profile with these settings:

VE_TSMCLI_NODE_NAME

Specify the node that connects the Data Protection for VMware command-line interface to the Tivoli Storage Manager server and the agent node (`VC1_VCLI1`).

VE_VCENTER_NODE_NAME

Specify the virtual node that represents a vCenter (`VC1`).

VE_DATACENTER_NAME

Specify the virtual node that maps to a data center. The correct syntax

is shown here:

datacenter_name::datacenter_node_name

- The datacenter_name value is case-sensitive.
- Make sure to set this parameter for each data center in your environment (VC1_DC1, VC1_DC2).
- The Data Protection for VMware vCenter plug-in does not support data centers with the same name in the vCenter.

VE_TSM_SERVER_NAME

Specify the hostname or IP of the Tivoli Storage Manager server.

VE_TSM_SERVER_PORT

Specify the port name to use for the Tivoli Storage Manager server. The default value is 1500.

An example profile with these settings is provided here:

```
VE_TSMCLI_NODE_NAME    VC1_VCLI1
VE_VCENTER_NODE_NAME   VC1
VE_DATACENTER_NAME     MyDatacenter::VC1_DC1
VE_DATACENTER_NAME     MyDatacenter::VC1_DC2
VE_TSM_SERVER_NAME     tsmserver.xyz.yourcompany.com
VE_TSM_SERVER_PORT     1500
```

See “Profile parameters” on page 64 for the complete list of available parameters, descriptions, and examples.

2. Verify that the IBM WebSphere Application Server (eWAS) service is running:

- **Windows** Click **Start > Control Panel > Administrative Tools > Services** and verify that the status of IBM WebSphere Application Server V7.0 - TSMVEplugin is Started.
- **Linux** Go to the scripts directory (/opt/tivoli/tsm/tdpvmware/common/scripts/) and issue this command:
./ewas status

These init scripts can also be used to stop and start the eWAS service:

```
/etc/init.d/ewas stop
/etc/init.d/ewas start
```

3. Verify that the client acceptor is running:

Windows Click **Start > Control Panel > Administrative Tools > Services** and verify that the status of Data Protection for VMware command-line interface is Started.

Linux Go to the scripts directory (/opt/tivoli/tsm/tdpvmware/common/scripts/) and issue this command:

```
./vmclid status
```

- If the daemon is running, proceed to Step 4.
- If the daemon is not running, issue this command to manually start the daemon:

```
/opt/tivoli/tsm/tdpvmware/common/scripts/vmcli --daemon
```

These init scripts can also be used to stop and start the daemon:

```
./vmclid stop
./vmclid start
```

4. Make sure the Tivoli Storage Manager Administrator has registered the node (to be used as the Data Protection for VMware command-line interface node) on the Tivoli Storage Manager server. Request the password for this node and set it in your pwd.txt file as shown in these examples:

- **Windows** Make sure that a space does not exist between the password (password1) and the greater-than sign (>).

```
echo password1> pwd.txt
```

or

```
echo password1>pwd.txt
```

```
vmcli -f set_password -I pwd.txt
```

- **Linux** You must issue the **vmcli -f set_password** command as tdpvmware user, and not as root.

```
echo password1 > pwd.txt
```

```
vmcli -f set_password -I pwd.txt
```

Important: When the Data Protection for VMware command-line interface node password is set, the password in the pwd.txt file is automatically changed to a 64 character random password. As a result, the original password is no longer valid. Additional information about this command and password behavior is available in “Set_password” on page 63.

5. Issue this vmcli command to verify that the Data Protection for VMware command-line interface recognizes the Tivoli Storage Manager node configuration:

```
vmcli -f inquire_config -t TSM
```

See “Inquire_config” on page 61 for details about this command.

Start the Data Protection for VMware vCenter plug-in

Complete these tasks after configuring the Data Protection for VMware vCenter plug-in.

1. Start the vSphere Client and log on to the vCenter. If the vSphere Client is already running, you must stop and restart it.
2. Go to the Home directory in the vSphere Client. Click the Data Protection for VMware vCenter plug-in icon in the Solutions and Applications panel.

Tip: If the icon is not shown, then the Data Protection for VMware vCenter plug-in was not registered or a connection error occurred.

- a. In the vSphere Client menu, go to **Plug-ins > Manage Plug-ins** to start the Plug-in Manager.
 - b. If you can locate the Data Protection for VMware vCenter plug-in and a connection error occurred, complete these tasks:
 - 1) Verify connectivity to the machine where the Data Protection for VMware vCenter plug-in is installed by issuing the ping command.
 - 2) Verify that IBM WebSphere Application Server V7.0 is running.
3. When the Data Protection for VMware vCenter plug-in starts for the first time, edit the Tivoli Storage Manager server settings by clicking **Configuration tab > Tivoli Storage Manager Server > Edit**.

The Data Protection for VMware vCenter plug-in is ready for backup and restore operations.

Related tasks

“Backing up virtual machine data to Tivoli Storage Manager”

Chapter 6, “Using the Data Protection for VMware command-line interface,” on page 57

Backing up virtual machine data to Tivoli Storage Manager

Back up your virtual machine data to Tivoli Storage Manager storage. You can run the task immediately or schedule it to run later.

Before proceeding, make sure the tasks described in “Configuring the Data Protection for VMware vCenter plug-in” on page 41 are completed successfully.

Restriction: Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters are not supported at this time.

Create a backup task for your virtual machine by following these steps:

1. Click the **Backup** tab or click **Define a backup task** to open the Managing backup schedules pane.
2. Click **Next** to begin the wizard. Follow the instructions in each page of the wizard and click **Next** to continue to the next page.
3. In the General page of the wizard, specify a name for the backup task you are creating in the **Backup Schedule Name** field. The name must not contain spaces. You can optionally add a description for the task.
4. In the Select what to back up page, expand the appropriate section of the tree and select the virtual machines that you want to back up.

Restriction: Backing up a virtual machine that contains a comma in its name is not supported.

5. In the Destination page, select the Tivoli Storage Manager data mover node that runs the backup operation. To avoid a backup failure, choose a data mover node that is not used by another process. Click **Next**.
6. In the Schedule page, you can choose to run the backup now or schedule the backup to run at a later time.
 - To begin the backup at the completion of this wizard, do the following steps:
 - a. Click **Run the backup now**.
 - b. Select an option in the **Backup type** field:
 - Click **Incremental backup** to back up the data that has changed since the last full backup. If a full backup has never been run on this virtual machine, a full backup is performed.
 - Click **Full backup** to back up an image of an entire virtual machine.
 - c. Click **Next** and skip to Step 9.
 - To schedule your task to run at a later time, do the following steps:
 - a. Click **Schedule for later**.
 - b. Select an option in the **Backup type** field:
 - Click **A full backup, followed by six incremental backups** to run a full backup weekly followed by six daily incremental backups. You must enter the name of the incremental backup schedule name and description.

- Click **Incremental backup** to back up the data that has changed since the last full backup. If a full backup has never been run on this virtual machine, a full backup is performed.
 - Click **Full backup** to back up an image of an entire virtual machine.
- c. Click **Next**.
7. If you chose to run a full backup followed by daily incremental backups, complete the Schedule repetition settings page:
 - a. Specify the date and time to run the first backup. The first full backup is scheduled to run at this date and time. The six incremental backups are scheduled to run on the remaining six days of the week and at the selected time.
 - b. If you want to include newly added or discovered virtual machines to future runs of this task, check **Newly added virtual machines are included in this backup task**: This check box has the following characteristics:
 - If you select all of the virtual machines on one ESX host and you select this option, the schedule backs up that ESX host. That is, all virtual machines on that host, present and future, are backed up.
 - If you select one or more virtual machines on an ESX host (but not all), and you select this option, then the schedule includes only the selected virtual machines and any future virtual machines that are added to the host. The remaining virtual machines on that host that are not selected are excluded.
 - If you do not select this option, only the virtual machines you selected are backed up, and no future virtual machines are included.
 8. If you chose to run a full backup or an incremental backup, complete the Schedule when the backups repeat page:
 - a. Specify the date and time to run the first backup.
 - b. Specify the interval that you want the backup to run.
 - c. If you want to include newly added or discovered virtual machines to future runs of this task, check **Newly added virtual machines are included in this backup task**. See Step 7b for details regarding this option.
 9. In the Ready to complete page, review your backup settings and click **Finish** to save your task. If you selected to run the backup now, the backup operation begins immediately when you click **Finish**.

After the backup task has completed, you can verify that the virtual machines have been backed up in the “Reporting” on page 53.

Back up options

The options you select for affect the way a virtual machine backup is created and stored.

Full backup

Click this option to back up an image of an entire virtual machine and send it to the Tivoli Storage Manager server.

Incremental backup

Click this option to back up only the data that has changed since the last full backup to Tivoli Storage Manager. If a full backup has never been run on this virtual machine, a full backup is performed.

Restoring a single virtual machine from Tivoli Storage Manager

Use this procedure to restore a single virtual machine from Tivoli Storage Manager to its original location or an alternate location.

Restriction: Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters are not supported at this time.

Restore a single virtual machine by completing the following steps:

1. Click the **Restore** tab or click **Initiate a restore** in the Getting Started window to open the Restore page.
2. In the Restore page, select the data center, ESX host, and virtual machine in the tree. A list of available backups is shown in the table.
3. Select the backup version to restore and click **Restore** to open the Restore Virtual Machine wizard.

Note: All new data stores and data centers must be created before starting the restore wizard.

4. Click **Next** to begin the wizard. Follow the instructions in each page of the wizard and click **Next** to continue to the next page. The Source step shows the details of the selected virtual machine backup.
5. In the Destination step, select the destination of the selected virtual machine. You can restore the virtual machine to the following locations:

Original Location

If the original virtual machine still exists, the wizard fields are populated with the original virtual machine information by default.

Important: If the original virtual machine still exists, you cannot restore a virtual machine to the original location. You must first remove the existing virtual machine before doing the restore operation. You can restore a virtual machine to its original location only if the virtual machine does not exist.

Alternate Location

Enter the following information to restore the virtual machine to an alternate location:

- a. Specify the name of the new virtual machine to which the backup is to be restored.
- b. Select the data center to use for the restore destination.
- c. Select the ESX host to use for the restore destination.

Restriction: When attempting to view and restore data originally backed up by a Tivoli Storage Manager Backup-Archive client version earlier than 6.3, the ESX host name is shown as UNKNOWN.

- d. Select the data store used for the restore destination.
- e. Select the Tivoli Storage Manager data mover node on the vStorage backup server that runs the restore operation. The default data mover node is the one that did the most recent backup.

After you make your selection, click **Next**.

6. The Summary step shows all the options you have selected in the Restore Virtual Machine wizard. Review your restore settings and click **Finish** to start the restore task.

After the restore task has started, you can monitor its progress, or verify that the restore task has completed in the “Reporting” on page 53.

Restoring multiple virtual machines from Tivoli Storage Manager

You can restore multiple virtual machines to the original data store and host or to an alternate data store and host from Tivoli Storage Manager.

Restriction: Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters are not supported at this time.

Restore multiple virtual machines by completing the following steps:

1. Click the **Restore** tab or click **Initiate a restore** in the Getting Started window to open the Restore page.
2. In the Restore page, select the data center, ESX host, and virtual machines in the tree. A list of available backups is shown in the table.
3. Select the backup versions to restore and click **Restore** to open the Restore Virtual Machine wizard.

Note: All new data stores and data centers must be created before starting the restore wizard.

4. Click **Next** to start the wizard. Follow the instructions in each page of the wizard and click **Next** to continue to the next page. The Source step shows the details of the selected virtual machine backups.
5. In the Destination step, select the destination of the selected virtual machine. You can restore the virtual machine to the following locations:

Original Location

If the original virtual machine still exists, the wizard fields are populated with the original virtual machine information by default.

Important: If the original virtual machine still exists, you cannot restore a virtual machine to the original location. You must first remove the existing virtual machine before doing the restore operation. You can restore a virtual machine to its original location only if the virtual machine does not exist.

Alternate Location

Enter the following information to restore the virtual machines to an alternate location:

- a. Select the data center to use for the restore destination.
- b. Select the ESX host to use for the restore destination.

Restriction: When attempting to view and restore data originally backed up by a Tivoli Storage Manager Backup-Archive client version earlier than 6.3, the ESX host name is shown as UNKNOWN.

- c. Select the data store used for the restore destination.

- d. Select the Tivoli Storage Manager data mover node on the vStorage backup server that runs the restore operation. The default data mover node is the one that did the most recent backup.

After you make your selection, click **Next**.

6. Review all the specified options as shown in the Summary step. Click **Finish** to start the restore operation.

After the restore task has started, you can monitor its progress, or verify that the restore task has completed in the “Reporting.”

Reporting

Use the Reports tab of the Data Protection for VMware vCenter plug-in to display progress information about your tasks and space usage information about your backups. This information can help you troubleshoot your backups.

Events

An event is a message that provides status on an operation such as a backup or restore.

To find out whether a backup or restore operation has completed or failed, click **Events** in the **View** field. The table in the Reports tab shows the backup tasks or restore tasks that have run and their status.

To see detailed information about an event, select an item from the table, and the details about the event appear in the information pane near the bottom of the tab. Click **Refresh** to update the information in the table.

Active Tasks

You can monitor the progress of a backup or restore task that is currently running. To display active tasks, click **Active Tasks** in the **View** field.

Tasks are initiated by activities that you perform in the Data Protection for VMware vCenter plug-in in real time, and by scheduled tasks that occur at a later time. If you keep the Active Tasks view open, the table and the details section of the tab are refreshed automatically. You can also manually refresh the table by clicking **Refresh**.

Information about all steps that have already been processed for an operation are displayed.

Data center Occupancy

Click **Data center Occupancy** in the **View** field to display details about the amount of space that is occupied by backups of Tivoli Storage Manager client nodes that are on the Tivoli Storage Manager server. The client nodes represent the VMware data center objects from which data is backed up to a Tivoli Storage Manager server. You can reduce a backup size by running incremental backups of a virtual machine.

This view is only available if a Tivoli Storage Manager server has been defined.

Refresh

Click **Refresh** to update the information in the table.

Troubleshooting

Solutions to Data Protection for VMware vCenter plug-in and Data Protection for VMware command-line interface issues are provided.

Data Protection for VMware vCenter plug-in backup or restore operation fails

Complete these tasks to resolve a backup or restore failure:

1. View these Tivoli Storage Manager backup-archive client log files to see if an error was generated:
 - `dsmerror.log`: All client messages.
 - `dsmwebcl.log`: All web client messages.

These log files are located in the directory you specify with the `DSM_LOG` environment variable or in the current working directory.

Tip: You can view error explanations in the Tivoli Storage Manager6.3 information center.

2. If neither of these files contain an error, run a Backup-Archive Client backup and restore operation to see if it fails.
3. If the backup-archive client operations complete successfully, run a Data Protection for VMware command-line interface “Backup” on page 57 and “Restore” on page 59 operation. Set the appropriate trace parameters (as described in “Profile parameters” on page 64) so you can view any errors that might be generated.

Data Protection for VMware command-line interface backup fails with `scSignOnAsAdmin: Error 53`

In this situation, a Data Protection for VMware command-line interface backup operation failed and this error was generated to the backup-archive client `dsmerror.log`:

```
scSignOnAsAdmin: Error 53 receiving SignOnAsAdminResp verb from server
```

Typically, this error results when the Data Protection for VMware command-line interface node name is different from its administrator name. Tivoli Storage Manager requires these two names to be the same.

For more information about node names, see “Register Tivoli Storage Manager nodes” on page 41.

Tivoli Storage Manager data mover nodes are not visible during a backup operation

Verify that the correct proxy node authority was granted on the Tivoli Storage Manager server as described in “Register Tivoli Storage Manager nodes” on page 41. If the correct authority exists, then the data center mapping specified by the `VE_DATACENTER_NAME` profile parameter is incorrect:

1. See “Profile parameters” on page 64 for a complete description and correct syntax of the `VE_DATACENTER_NAME` parameter.

2. (Optional) See “Configuring the Data Protection for VMware command-line interface” on page 46 if you must update your configuration.

The `inquire_detail` command failed with Return Code 53

In this situation, the `vmcli -f inquire_detail` command failed and this error was generated to your log file:

```
ANS1033E (RC-53) An invalid TCP/IP address was specified.
```

This error occurs when a node name does not match its administrator name. This issue can happen when you rename a node but do not rename its administrator. The solution is to either rename the administrator to match the new node name or register a new administrator for the new node.

The commands in these examples are issued from the Tivoli Storage Manager Administrative Command Line:

- Rename the administrator at the same time you rename the node:

```
rename node <current_node_name> <new_node_name>  
rename admin <current_admin_name> <new_node_name>
```

For example:

```
rename node DC_VC5 DC_WIN2K8_X64  
rename admin DC_VC5 DC_WIN2K8_X64
```

As a result, the new administrator name matches the new node name.

- Register the administrator directly after renaming the node:

```
rename node <current_node_name> <new_node_name>  
register admin <new_admin_name> <password>
```

For example:

```
rename node DC_VC5 DC_WIN2K8_X64  
register admin DC_WIN2K8_X64 DC_WIN2K8_X64PWD
```

As a result, the new administrator name matches the new node name.

Invalid sign on and invalid password errors received

In this situation, Data Protection for VMware command-line interface or Data Protection for VMware vCenter plug-in operations failed and your log file contained one (or more) of these errors:

```
GVM1170E: A VMCLI command failed.
```

```
ANR2177I FRSV123015.TSMCLI has 1 invalid sign-on attempts. The limit is 5.
```

```
ANR0424W Session 125713 for node FRSV128215.TSMCLI (TSM4VE)  
refused - invalid password submitted.
```

This issue can occur when a Data Protection for VMware command-line interface node set password attempt was made by using a password that was previously set.

| When the Data Protection for VMware command-line interface node password is
| set, the password is registered and then automatically changed. As a result, you do
| not know this changed password. If you must change the Data Protection for
| VMware command-line interface node password after it is already set, the Tivoli
| Storage Manager Administrator must use the update node command to reset this
| node on the Tivoli Storage Manager server. When the node has been reset, set the
| Data Protection for VMware command-line interface node password again and
| specify the new password.

| The Data Protection for VMware command-line interface node password is set by
| either of these methods:

- | • Data Protection for VMware command-line interface: `vmcli -f set_password`
| command
- | • Data Protection for VMware vCenter plug-in: Configuration tab -> Set Node
| Password page

| **Session timeout**

| The Tivoli Storage Manager server COMMTIMEOUT option affects the duration of
| the Data Protection for VMware session. If the processing time of the Data
| Protection for VMware operation exceeds this value, the server ends the session
| with Data Protection for VMware. Therefore, if you are sure that no error has
| occurred during a Data Protection for VMware operation and the
| COMMTIMEOUT value has been reached, increase the value. Likewise, if an error
| occurred but Data Protection for VMware did not report the error in a timely
| manner, then decrease the value for better real-time reporting.

Chapter 6. Using the Data Protection for VMware command-line interface

Use the Data Protection for VMware command-line interface to back up, restore, or view configuration information.

Complete the tasks described in “Configuring the Data Protection for VMware command-line interface” on page 46 before issuing the vmcli commands.

The Data Protection for VMware command-line interface provides these commands:

“Backup”

Initiate full and incremental backups of your virtual machines.

“Restore” on page 59

Restore backups of your virtual machines.

“Inquire_config” on page 61

View configuration information about the backup database.

“Inquire_detail” on page 62

View configuration information about the backup environment.

“Set_password” on page 63

Set the password for the Data Protection for VMware command-line interface node name.

Backup

Use this vmcli command to start full and incremental backups of your virtual machines.

Syntax

The **vmcli -f backup** command uses this syntax:

```
vmcli -f backup -t backupType -I backupObjectListFile -d datacenternodename -o datamovernodename [--name taskName] [--description descriptionInFile.txt] [-s tmsserverhostname][ -n vctrclinodename] [-p tmsserverport]
```

Linux

You must issue the **vmcli -f backup** command as `tdpvmmware` user, and not as root.

Parameters

Before issuing a **vmcli -f backup** command, issue the **vmcli -f inquire_config** command to verify that your configuration is correct. Also, use the information from the **vmcli -f inquire_config** command output as a guide for setting your backup parameters.

| **-t** *backupType*
| Specify the type of backup to complete. You can choose from one of the
| following types:

| **TSM_INCR**

| Creates an incremental backup of the specified backup object. This is
| the default.

| **TSM_FULL**

| Creates a full image backup of the specified backup objects.

| **-I** *backupObjectListFile*

| Specify the file that contains the list of objects to backup. The
| *backupObjectListFile* uses the following keyword:

| **vmname**

| Specify the name of the VM to back up. You can specify this keyword
| for each VM you want to back up. For example:

```
vmname: vm1  
vmname: vm2
```

| **Restriction:** When specifying the name of a VM using the *vmname*
| keyword in the *backupObjectListFile*, Data Protection for VMware does
| not differentiate between a colon (:) used as a keyword separator or a
| colon used in the VM name. Therefore, use caution when specifying
| keyword values. In addition, backing up a VM that contains a comma
| in its name is not supported. In addition, Data Protection for VMware
| support for virtual machine backup and restore operations is limited to
| virtual machine names and data center names that contain English 7
| bit ASCII characters only. Virtual machine names and data center
| names that use other language characters are not supported at this
| time.

| **-d** *datacenternodename*

| Specify the data center node name.

| **-o** *datamovernodename*

| Specify the Tivoli Storage Manager data mover node name. This name is the
| node name for the Tivoli Storage Manager Backup-Archive Client that is
| installed on the vStorage Backup Server. This node performs the data
| movement.

| **[--name** *taskName*]

| Specify the string that identifies the backup task.

| **[--description** *descriptionInFile.txt*]

| Specify the name of the text file that contains a description of the backup task.

| **[-s** *tmsserverhostname*]

| Specify the host name or IP address of the Tivoli Storage Manager server. If
| this parameter is not specified, the value in the profile is used.

| **[-n** *vctrclinodename*]

| Specify the Data Protection for VMware command-line interface node name.
| This is the node that connects the Data Protection for VMware command-line
| interface to the Tivoli Storage Manager server and the Tivoli Storage Manager
| data mover node. If this parameter is not specified, the value in the profile is
| used.

[-p *tmsserverport*]

Specify the port of the Tivoli Storage Manager server.

- If this parameter is not specified in the Data Protection for VMware command-line interface and not specified in the profile, the default port (1500) is used.
- If this parameter is not specified in the Data Protection for VMware command-line interface, but is specified in the profile, the value in the profile is used.

Restore

Use this `vmcli` command to restore backups of your virtual machines.

Syntax

The `vmcli -f restore` command uses this syntax:

```
vmcli -f restore -I restoreObjectListFile -d datacenternodename -o datamovernodename [-s tmsserverhostname] [-n vctrclinodename] [-p tmsserverport]
```

Linux

You must issue the `vmcli -f restore` command as `tdpvmware` user, and not as `root`.

Parameters

-I *restoreObjectListFile*

Specify the file that contains the list of VMs to restore. Each line can contain only one VM identifier.

The *restoreObjectListFile* uses the following keywords:

backupid

Specify the Tivoli Storage Manager Object ID for a specific VM backup. Locate the Object ID using the `vmcli -f inquire_detail` command. This keyword is required for a restore operation.

vmname

Specify the name of the VM that was originally backed up. If this keyword is not specified, the name `vmname` is used for the restore.

Restriction: When specifying a keyword in the *restoreObjectListFile*, Data Protection for VMware does not differentiate between a colon (:) used as a keyword separator or a colon used in a keyword value. Therefore, use caution when specifying keyword values. In addition, Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters are not supported at this time.

vmname

(Second entry) Specify the name that you want the restored VM to be named. Existing VMs are not overwritten. Therefore, either rename the VM (using this keyword) or delete the original VM before issuing the `vmcli -f restore` command.

| **newdatacentername**

| When you want the restore destination to be a different data center,
| specify the name of that data center with this keyword.

| **newesxhostname**

| When you want the restore destination to be a different ESX host,
| specify the name of that ESX host with this keyword.

| **newdatastoreurl**

| Specify the name (not the URL) of the data store where the VM is to be
| restored. For example, a data store name such as datastore1 is
| supported. A data store URL such as sanfs://vmfs_uuid:4d90pa2d-
| e9ju45ab-065d-00101a7f1a1d/ is not supported.

| An example *restoreObjectListFile* is provided here:

```
# restore of VM "678912345" named "vmName6" to new vmname "vm6newName" to datacenter  
"DataCenter2" to ESX esxhostname:esxHost1Name to new datastore "datastore2"  
backupid:678912345 vmname:vmName6::vmname:vm6newName newdatacentername:DataCenter2  
newesxhostname:esxHost1Name newdatastoreurl:datastore2
```

| Each restore specification must be on a single line. However, for the sake of
| page formatting, the restore specification in this example is on multiple lines.

| **Tip:** To make sure that correct information is specified in the
| *restoreObjectListFile*, you can issue the **inquire_detail** command.
| "Inquire_detail" on page 62 provides current configuration information about
| the backup environment.

| **-d** *datacenternodename*

| Specify the data center node name.

| **-o** *datamovernodename*

| Specify the Tivoli Storage Manager data mover node name. This is the node
| name for the Tivoli Storage Manager Backup-Archive Client that is installed on
| the vStorage Backup Server. This node performs the data movement.

| **[-s** *tsmserverhostname*]

| Specify the host name or IP address of the Tivoli Storage Manager server. If
| this parameter is not specified, the value in the profile is used.

| **[-n** *vctrclinodename*]

| Specify the Data Protection for VMware command-line interface node name.
| This name is the node that connects the Data Protection for VMware
| command-line interface to the Tivoli Storage Manager server and the Tivoli
| Storage Manager data mover node. If this parameter is not specified, the value
| in the profile is used.

| **[-p** *tsmserverport*]

| Specify the port of the Tivoli Storage Manager server.

- If this parameter is not specified in the Data Protection for VMware command-line interface and not specified in the profile, the default port (1500) is used.
- If this parameter is not specified in the Data Protection for VMware command-line interface but is specified in the profile, the value in the profile is used.

Inquire_config

Use this `vmcli` command to view configuration information about the backup database that is associated with Data Protection for VMware.

Syntax

The `vmcli -f inquire_config` command uses this syntax:

```
vmcli -f inquire_config [-v vcenternodename] [-s tmsserverhostname] [-n vctrclinodename] [-p tmsserverport]
```

Linux

You must issue the `vmcli -f inquire_config` command as `tdpvmware` user, and not as `root`.

Parameters

[-v *vcenternodename*]

Specify the virtual node that represents that represents a vCenter. If this parameter is not specified in the Data Protection for VMware command-line interface, the value in the profile is used.

[-s *tmsserverhostname*]

Specify the host name or IP address of the Tivoli Storage Manager server. If this parameter is not specified, the value in the profile is used.

[-n *vctrclinodename*]

Specify the Data Protection for VMware command-line interface node name. This name is the node that connects the Data Protection for VMware command-line interface to the Tivoli Storage Manager server and the Tivoli Storage Manager data mover node. If this parameter is not specified, the value in the profile is used.

[-p *tmsserverport*]

Specify the port of the Tivoli Storage Manager server.

- If this parameter is not specified in the Data Protection for VMware command-line interface and not specified in the profile, the default port (1500) is used.
- If this parameter is not specified in the Data Protection for VMware command-line interface but is specified in the profile, the value in the profile is used.

Example

The parameter values in this output for the `vmcli -f inquire_config -s TSM` command show that the Data Protection for VMware command-line interface recognizes the Tivoli Storage Manager node configuration. As a result, the configuration is correct:

```

#TASK 858 inquire_config 2011082223402852
#PARAM INSTALLED=FCM
#PARAM INSTALLED=TSM
#RUN 860 2011082223402852
#LANG en_US
#PARAM BACKEND=TSM
#PARAM OPERATION_TYPE 4
#PHASE_COUNT 4
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=TEMPLAR.STORAGE.USCA.IBM.COM
#PARAM TSMCLI_NODE_NAME=DPM02_TSMCLI
#PARAM VCENTER_NODE_NAME=DPM02_VCENTER
#PARAM DATACENTER_NODE_NAME=
#PARAM OFFLOAD_HOST_NAME=
#PARAM TSM_OPTFILE=/tmp/T4VE_DtYWYf
#PARAM INPUT_FILE=
#PARAM TRACEFILE=
#PARAM TRACEFLAGS=
#PHASE INITIALIZE
#PHASE INQUIRE_DATACENTER_NODES
#CHILD datacenternode:DATACENTER1::DPM02_DC1
#PARENT vcenternode:DPM02_VCENTER
#PHASE INQUIRE_PROXY_NODES
#CHILD targetnode:DPM02_DC1
#PARENT peernode:DPM02_DC1_DM
#CHILD hladdress:untoro.storage.usca.ibm.com
#PARENT peernode:DPM02_DC1_DM
#CHILD lladdress:62790
#PARENT peernode:DPM02_DC1_DM
#PARAM STATUS=sucesss
#END RUN 860 20110822234030547
#END TASK 858
#INFO FMM16014I The return code is 0.
#END

```

Inquire_detail

Use this `vmcli` command to view configuration information about the backup environment associated with Data Protection for VMware.

Syntax

The `vmcli -f inquire_detail` command uses this syntax:

```
vmcli -f inquire_detail -d datacenternodename [-a] [-n vctrclinodename] [-p tsmsserverport] [-q vmfs | vmsingle (-I inputfile)] [-s tsmsserverhostname]
```

Linux

You must issue the `vmcli -f inquire_detail` command as `tdpvmware` user, and not as root.

Parameters

-d *datacenternodename*
Specify the data center node name.

[-a]
Specify to show only the active backups on the Tivoli Storage Manager server.

[-n *vctrclinodename*]
Specify the Data Protection for VMware command-line interface node name.
This name is the node that connects the Data Protection for VMware

command-line interface to the Tivoli Storage Manager server and the Tivoli Storage Manager data mover node. If this parameter is not specified, the value in the profile is used.

[-p *tmsserverport*]

Specify the port of the Tivoli Storage Manager server.

- If this parameter is not specified in the Data Protection for VMware command-line interface and not specified in the profile, the default port (1500) is used.
- If this parameter is not specified in the Data Protection for VMware command-line interface but is specified in the profile, the value in the profile is used.

[-q *vmfs* | *vmsingle* (-I *inputfile*)]

Specify **vmfs** to query all VMware Virtual Machine File Systems (VMFS).

Specify **vmsingle** to query all VM Guests identified in the *inputfile*. When **q** is specified without the **vmfs** or **vmsingle** value, the default value is to query all file spaces.

[-s *tmsserverhostname*]

Specify the host name or IP address of the Tivoli Storage Manager server. If this parameter is not specified, the value in the profile is used.

Set_password

Use this `vmcli` command to set the password for the Data Protection for VMware command-line interface node name.

Syntax

The `vmcli -f set_password` command uses this syntax:

```
vmcli -f set_password -I passwordfile
```

Linux

You must issue the `vmcli -f set_password` command as `tdpvmware` user, and not as `root`.

Parameters

-I *passwordfile*

Specify the password for the Data Protection for VMware command-line interface node name. The *passwordfile* must contain only the password and no other content. The *passwordfile* is deleted by the Data Protection for VMware command-line interface when the `vmcli -f set_password` command is issued.

Important: When the Data Protection for VMware command-line interface node password is set, the password is registered and then automatically changed to a 64 character random password. As a result, the original password is no longer valid and you do not know the random password. If you must change the Data Protection for VMware command-line interface node password after it is already set, the Tivoli Storage Manager Administrator must use the `update node` command to reset this node on the Tivoli Storage Manager server. When the node has been reset, set the Data Protection for VMware command-line interface node password again and specify the new password.

Example

Windows When creating the password file using the echo command, make sure that a space does not exist between the password (password1) and the greater-than sign (>). For example:

```
echo password1> pwd.txt
```

or

```
echo password1>pwd.txt
```

This example sets the password (password1) in file pwd.txt:

```
vmcli -f set_password -I pwd.txt
```

Linux Create the password file (pwd.txt) using the echo command:

```
echo password1 > pwd.txt
```

This example sets the password (password1) in file pwd.txt:

```
vmcli -f set_password -I pwd.txt
```

Profile parameters

Use the Data Protection for VMware command-line interface profile to configure settings for backup and restore tasks in your environment.

The profile is located in this directory on the system where the Data Protection for VMware vCenter plug-in is installed:

Linux /home/username/tdpvmware/config

Windows C:\Program Files\Common Files\Tivoli\TDPVMware\VMwarePlugin\scripts

DERBY_HOME <path to Derby database>

This parameter specifies the location of the Derby database that is used by the Data Protection for VMware command-line interface.

Example:

Linux

```
DERBY_HOME /opt/tivoli/tsm/tdpvmware/derby
```

Windows

```
DERBY_HOME C:\Program Files\Common Files\Tivoli\TDPVMware\VMwarePlugin\derby
```

VE_DATACENTER_NAME <data_center_name::DATA_CENTER_NODE_NAME>

Specify the data center (data center name) with a value that is case-sensitive and that matches the data center name used in the vCenter. Specify the virtual node (DATA_CENTER_NODE_NAME) that maps to the data center. If the vCenter manages several data centers, you can specify this parameter for each data center. However, the Data Protection for VMware vCenter plug-in does not support data centers with the same name in the vCenter.

Example:

```
VE_DATACENTER_NAME DataCenter1::Fin_Datacenter1
VE_DATACENTER_NAME DataCenter2::Fin_Datacenter2
```

Restriction: Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters are not supported at this time.

After a data center name is created and associated with a Tivoli Storage Manager node, be aware of these restrictions:

- Do not change the data center name in the vCenter without also creating a Tivoli Storage Manager node name and associating it with the new data center name.
- Do not change the data center name and the profile without also changing the Tivoli Storage Manager node name.
- Do not create a data center mapping value in the profile with a previously used Tivoli Storage Manager node.

When the data center name in the vCenter has changed, you must complete these steps before attempting any operations:

1. Register a data center node for the new data center name.
2. Grant proxy authority to the new data center node to perform tasks on behalf of the vCenter node.
3. Update the profile with the new data center mapping.
4. Grant proxy authority to the Tivoli Storage Manager data mover nodes to perform tasks on behalf of the new data center node.
5. Remove any entry from the profile that used the previous data center node or vCenter node name.

VE_TRACE_FILE *<path and name of trace file>*

Specify the full path and name of the file to be used to contain trace information. Activate tracing only when instructed to do so by IBM Software Support.

VE_TRACE_FLAGS *<flags>*

Specify one or more trace flags. Multiple trace flags are separated with a space. Activate tracing only when instructed to do so by IBM Software Support.

VE_TSMCLI_NODE_NAME *<Data Protection for VMware command-line interface node>*

Specify the Data Protection for VMware command-line interface node. This node connects the Data Protection for VMware command-line interface to the Tivoli Storage Manager server and Tivoli Storage Manager data mover node.

Example:

```
VE_TSMCLI_NODE_NAME VC1_VCLI1
```

VE_TSM_SERVER_NAME *<server host name or IP address>*

Specify the host name or IP address of the Tivoli Storage Manager server used for backup operations. There is no default value.

Example:

```
VE_TSM_SERVER_NAME tsmserver.xyz.yourcompany.com
```

VE_TSM_SERVER_PORT <port name>

Specify the port name to use for the Tivoli Storage Manager server. The default value is 1500.

Example:

```
VE_TSM_SERVER_PORT 1500
```

VE_VCENTER_NODE_NAME <vCenter node>

Specify the vCenter node. This virtual node represents a vCenter.

Example:

```
VE_VCENTER_NODE_NAME VC1
```

VMCLI_DB_BACKUP *NO AT[day[, day[,....]]] time TO backup location*

This parameter controls the backup of the Derby database containing the metadata of the Data Protection for VMware command-line interface. Specify one of these values:

NO This option does not perform a backup of the Derby database.

AT [day[, day[,....]]] *time_in _24_H*

This option creates a backup on the specified day or days at the specified time, which is triggered by the scheduler. If the day value is not specified, a daily backup is created. Specify one of these values: MON, TUE, WED, THU, FRI, SAT, SUN.

You can separate these values by a comma or a blank space.

AFTER_BACKUP

This option creates a backup of the Derby database after each Data Protection for VMware backup operation.

The default location for the backups of the Derby database is *install_dir/derby_backups*. Specify *TO path* to set a custom path.

Example:

```
VMCLI_DB_BACKUP AT 00:00
```

VMCLI_DB_BACKUP_VERSIONS <number>

Specify the maximum number of backup generations that are maintained for the Derby database, before the oldest version is overwritten by a new version. This parameter applies only to the backups of the Derby database containing metadata. It has no effect on the number of backup generations that are maintained for the backups of a vSphere environment. The default value is 3.

Example:

```
VMCLI_DB_BACKUP_VERSIONS 3
```

VMCLI_DB_HOST <Derby database local host name>

Specify the local host name of the Derby database. You can specify the host name (localhost) or the IP address (0.0.0.0).

Example:

```
VMCLI_DB_HOST localhost
```

VMCLI_DB_NAME <Derby database name>

Specify the name of the Derby database. The default value is VMCLIDB.

Example:

```
VMCLI_DB_NAME VMCLIDB
```

VMCLI_DB_PORT <Derby database port number>

Specify the Derby database port on which the Data Protection for VMware command-line interface starts and connects to the database. The default value is 1527. If this port is in use by another application, specify a different port.

Example:

```
VMCLI_DB_PORT 1527
```

VMCLI_GRACE_PERIOD <seconds>

When a backup is no longer available on Tivoli Storage Manager, the backup is marked for deletion as defined by a deletion date. However, before the backup is deleted, a grace period exists. Use this parameter to specify the grace period (length of time) between the deletion date and the date the backup is deleted from the Derby database. The default value is 2592000 seconds (30 days).

Example:

```
VMCLI_GRACE_PERIOD 1296000
```

VMCLI_LOG_DIR <path of log file>

Specify the absolute location or the relative location of the installation directory where the Data Protection for VMware command-line interface writes its log files. The default value is logs. If the default value logs is used, then all logs (and trace information) are written to these locations:

```
Linux /opt/tivoli/tsm/tdpvmware/common/logs
```

```
Windows C:\Program Files\Common Files\Tivoli\TDPVMware\logs
```

Example:

```
VMCLI_LOG_DIR logs
```

VMCLI_RECON_INTERVAL_TSM <seconds>

This parameter specifies the interval between *reconciliation* operations on the Derby database with Data Protection for VMware. Reconciliation operations delete metadata for backups that are no longer available. This action ensures the Derby database remains synchronized with the Data Protection for VMware repository. The default value is 1200 seconds.

Example:

```
VMCLI_RECON_INTERVAL_TSM 1200
```

VMCLI_RESTORE_TASK_EXPIRATION_TIME <seconds>

Specify the time that a Data Protection for VMware command-line interface restore task is stored in the Derby database. The default value is 2592000 seconds (30 days).

Example:

```
VMCLI_RESTORE_TASK_EXPIRATION_TIME 2592000
```

VMCLI_SCHEDULER_INTERVAL <seconds>

Specify the interval, in seconds, between scheduler checks for scheduled tasks due for execution. The default value is 1 second.

Example:

```
VMCLI_SCHEDULER_INTERVAL 60
```

VMCLI_TASK_EXPIRATION_TIME <seconds>

This parameter specifies the time that a task is stored in the Data Protection for VMware command-line interface Derby database. This parameter applies only to the **inquire_config** command. The default value is 864000 seconds (10 days).

Example:

```
VMCLI_TASK_EXPIRATION_TIME 864000
```

VMCLI_TRACE YES|NO

Specify that tracing files are activated. Activate tracing only when instructed to do so by IBM Software Support.

Profile example

An example profile is provided here:

```
DERBY_HOME                C:\test\Tivoli\TDPVMware\VMwarePlugin\derby
VE_DATACENTER_NAME        MyDataCenter::VC1_DC1
VE_DATACENTER_NAME        MyDataCenter::VC1_DC2
VE_TSMCLI_NODE_NAME       VC1_VCLI1
VE_TSM_SERVER_NAME        tsmserver.xyz.yourcompany.com
VE_TSM_SERVER_PORT        1500
VE_VCENTER_NODE_NAME      VC1
VMCLI_DB_BACKUP           AT 00:00
VMCLI_DB_BACKUP_VERSIONS  3
VMCLI_DB_HOST             localhost
VMCLI_DB_NAME             VMCLIDB
VMCLI_DB_PORT             1527
VMCLI_GRACE_PERIOD        2592000
VMCLI_LOG_DIR             logs
VMCLI_RECON_INTERVAL_TSM  1200
VMCLI_RESTORE_TASK_EXPIRATION_TIME 2592000
VMCLI_SCHEDULER_INTERVAL  60
VMCLI_TASK_EXPIRATION_TIME 864000
```

Chapter 7. Using the Data Protection for VMware Recovery Agent

The Data Protection for VMware Recovery Agent can mount snapshots to enable file level restores and can perform instant restores of volumes.

Backing up VMware virtual machine data

Windows

You can back up a VMware guest by using a full VMware backup.

The full VM backup method can back up any VMware guests that are supported by VMware. If you have Data Protection for VMware and are using the VStore API backup method, you can recover individual files, have a single source full VM backup, run incremental block level backups, and perform instant restore of volumes. For example:

```
backup vm my_vm -mode=incremental
```

If the VMware client is using Tivoli Storage Manager version 6.2.3 or later, you can use the VStorage API to run a full backup. Data Protection for VMware is not a requirement for this feature. The VStore API offers the following features:

- Backing up VMware guests do not use as a temporary directory. Therefore, the Vstore API does not use as much storage on the off-host backup server.
- If you have Data Protection for VMware, you can restore individual files.
- If you have Data Protection for VMware, you can run incremental backups after running an initial full backup.
- If a backup was created using the Vstore API, you can restore the backup without using the VMware Converter tool.

Restriction: Data Protection for VMware support for virtual machine backup and restore operations is limited to virtual machine names and data center names that contain English 7 bit ASCII characters only. Virtual machine names and data center names that use other language characters is not supported at this time.

Scenario: Backing up your virtual machines

Windows

Schedule full and incremental backups to protect your virtual machine. This scenario guides you through the recommended settings for your VMware guests and the Tivoli Storage Manager client to implement VMware backups.

In this scenario, you define a full VMware backup of the guests that runs once a week and a daily incremental backup of the same VMware guests. This configuration ensures that there are frequent backups of the VMware guests and reduces the size of each backup.

Tip: There is no limit to how many full and incremental backups you can take. However, if you do not run a full backup regularly, the size of incremental backups can increase. This scenario ensures that the incremental backups do not get too large.

1. Prepare the off-host server for backups. See “Preparing the environment for VMware backup processing.”
2. Define a full VM backup for each of the VMware guests. See “Running full virtual machine backups” on page 73.
3. Separate the pool of VMware guests into groups to reduce backup time. The backup time is shorter because each group is backed up by a separate instance on the Tivoli Storage Manager client scheduler, and they are running in parallel on the off-host backup server.
4. Schedule the full VM backup that runs once a week.
5. Configure daily backups using “Running incremental virtual machine backups using the VStore API” on page 74.
6. Schedule the incremental backups to run daily.
7. Specify compression and deduplication to reduce the backup size of the VMware backups using the following steps:
 - a. From the VMware guests, open the Tivoli Storage Manager client `dsm.opt` file.
 - b. Enable compression by adding the following option to the client `dsm.opt` file: `compression yes`.

Tip: You can only enable compression if you are using client deduplication and if deduplication has been enabled for the storage pool.

- c. Enable deduplication by adding the following option to the client `dsm.opt` file: `deduplication yes`.

Preparing the environment for VMware backup processing

Use the following steps to prepare the VMware environment to be backed up. The vStorage backup server can run either a Windows or Linux client.

1. Configure your storage environment so that it can be backed using the following steps:
 - a. Configure your storage environment so that the vStorage backup server can access to the storage volumes that are in your ESX server farm.

Tip: By making the storage volumes visible, the backup uses the SAN transfer path during backup operations.
 - b. If you are using network attached storage (NAS) or direct-attach storage, ensure that the vStorage backup server is accessing the volumes with a network-based transport.
 - c. Optional: If you will be running a direct SAN backup, zone the SAN and configure the disk subsystem host mappings so that all VMware ESX servers and the vStorage backup server access the same disk volumes.
2. Configure the vStorage backup server using the following steps:
 - a.  Set and export the `LD_LIBRARY_PATH` environment variable to point to the client installation directory. For example, `export LD_LIBRARY_PATH=/opt/tivoli/tsm/client/ba/bin`. For convenience, each account that uses backup-archive client commands (for example, `dsmc`, `dsmcad`, or `dsmj`) should add this directory to their path.

- b. **Windows** When the backup-archive client runs on a vStorage backup server, this client configuration is referred to as the Tivoli Storage Manager data mover node (previously called the backup proxy server). If you configure the Tivoli Storage Manager data mover node to directly access the storage volumes, turn off automatic drive letter assignment. If it is not turned off, the client on the data mover node might corrupt the Raw Data Mapping (RDM) of the virtual disks and backups will fail.

A Tivoli Storage Manager data mover node typically uses the SAN to back up and restore data. If the Tivoli Storage Manager data mover node is on a Windows Server 2008 or Windows Server 2008 R2, the Windows SAN policy must be changed to `OnlineAll`. Setting the SAN policy to `OnlineAll` enables the operating system to bring new disks online, and make them available to the backup server. The policy must be set to `OnlineAll` to allow restore operations to use the `vmvstortransport hotadd` or `san transport` to restore data. Run **diskpart.exe** and enter the following commands:

```
diskpart
automount disable
automount scrub
san policy OnlineAll
exit
```

- c. **Windows** On Windows systems: Install the Tivoli Storage Manager client on the vStorage backup server. At the custom setup page of the installation wizard, select **VMware Backup Tools**.
3. Modify the Tivoli Storage Manager server using the following steps:
- a. Access the administrative command line of Tivoli Storage Manager client.
 - b. From the Tivoli Storage Manager client of the vStorage backup server, enter the following command to register the node:

```
register node my_server_name my_password
```

Where *my_server_name* is the full computer name of the vStorage backup server and *my_password* is the password to access it.

Tip: **Windows** You can get the server full computer name by right-clicking on **My Computer**. Click the Computer Name tab and look at the name listed next to **Full computer name**.

- c. **Windows** From the Tivoli Storage Manager client of the vStorage backup server, enter the following command to register the node:

```
register node my_guest_name my_password
```

where *my_guest_name* is the full name of the VMware guest that you are backing up. Repeat Step 3c for each VMware guest that you are backing up.

- d. Grant proxy access for each of the VMware guests by entering the following command:

```
grant proxynode agent=my_server_name target=my_guest_name
```

When the command completes, the following message is returned:

```
GRANT PROXYNODE: success.
Node my_server_name is granted proxy authority to node my_guest_name.
```

Repeat Step 3d for each of the VMware guests that you are backing up.

Virtual machine backup schedule examples

The Tivoli Storage Manager Backup-Archive Client provides a command-line interface and a graphical user interface. You can use either interface to perform vStorage virtual machine backups. You can run both full and incremental backups.

You can use the following client schedule options:

- Specify appropriate virtual machine schedule options.
 - mode=full or incremental. The full backup of a set of virtual machines requires a different schedule than incremental backup of the same set of virtual machines.
 - asnode=datacenternodename

Here are examples of two schedules that perform a weekly full backup and daily incremental backups for two ESX servers. Each dsmc instance processes an ESX server. The virtual node name datacenter1 represents a VMware data center. The options file for vmnode1 contains VMCH esx1, The options file for vmnode2 contains VMCH esx2.

- Schedule a weekly full backup of all virtual machines on ESX host esx1 and esx2. Use vmnode1 and vmnode2 instances and datacenter1 as the virtual node name that maps to data center.

```
define schedule vmdomain vmschedfullesx1 type=client action=backup
subaction=vm options='-asnodename=datacenter1 -mode=full'
startdate=mm/dd/yyyy starttime=hh:mm
perunits=weeks dayofweek=saturday
```

```
define association vmdomain vmschedfullesx1 vmnode1, vmnode2
```

- Schedule daily incremental backups of all virtual machines on ESX host esx1 and esx2. Use vmnode1 and vmnode2 instances and datacenter1 Use vmnode1 and vmnode2 instances and datacenter1 as the virtual node name that maps to data center.

```
define schedule vmdomain vmschedincesx1 type=client action=backup
subaction=vm options='-asnodename=datacenter1 -mode=inc'
startdate=mm/dd/yyyy starttime=hh:mm
schedstyle=enhanced perunits=weeks dayofweek==sunday, monday,
tuesday, wednesday, thursday, friday
```

```
define association vmdomain vmschedincesx1 vmnode1, vmnode2
```

Security considerations

The Tivoli Storage Manager Backup-Archive Client can back up a number of virtual machines under the same node name. Each virtual machine is a separate file space. The **dsmc set access** command can be used to control which virtual machines can be restored by which Mount or instant restore node name.

Two typical deployments are:

Mount installed on an off-host machine

The Tivoli Storage Manager administrator or help desk operator is responsible for mounting a snapshot and exporting it to the appropriate virtual machine. The **set access** command is issued from the backup-archive client node that owns the virtual machines to authorize the mount node to all virtual machines. For example,

```
set access backup * mountnodename
```

Mount and instant restore installed in-guest

The virtual machine user is responsible for restoring the data. To authorize

the mount or instant restore node to a specific virtual machine, issue the **set access** command from the backup-archive client node that owns the virtual machines. For example:

```
set access backup "{\VMFULL-vmdisplayname}\*\*" * mountnodename
```

The **set access** command does not restrict the ability to see what virtual machines have been backed up. However, it does restrict the ability to restore a virtual machine.

Running full virtual machine backups

Windows

A full virtual machine backup is a backup of an entire export of a virtual machine snapshot. It is like a Tivoli Storage Manager image backup.

The off-host VMware virtual machine must be configured as described in “Preparing the environment for VMware backup processing” on page 70.

1. Configure the backup Tivoli Storage Manager client on the off-host backup server using **Client Preferences** editor:
 - a. From the welcome page of the Tivoli Storage Manager client, click **Edit > Client Preferences**.
 - b. Click the notebook tab **VM Backup**.
 - c. Select **VMWare Full VM**.
 - d. In the **Domain Backup Types** list, select **Domain Full VM**.
 - e. In the **Host** field, enter either the host name of the vCenter or the host name of the ESX server. Enter the user ID and password information.

Note: Enter the Virtual Center host name. If you specify the Virtual Center host name, you can back up virtual machines from any of the VMware servers that are managed by the Virtual Center.

- f. In **VM Full Backup Type** section, select **VStorage**.
 - g. Click **OK** to save your changes.
2. Verify that your system is configured correctly by running a backup of one of the virtual machines using the following steps:

- a. At the command line of the off-host backup server, enter the command:

```
dsmc backup vm my_vm_name
```

where *my_vm_name* is the name of your virtual machine as seen in the VMware vSphere client.

- b. When the command ends, verify that it completed without errors. The following message indicates that the command completed successfully:

```
Backup VM command complete
Total number of virtual machines backed up successfully: 1
virtual machine vmname backed up to nodename NODE
Total number of virtual machines failed: 0
Total number of virtual machines processed: 1
```

3. Verify that you can restore the virtual machines files using the following steps:
 - a. At the command-line interface of the off-host backup server, enter the command:

```
dsmc restore vm my_vm_name -vmbackuptype=fullvm
```

- b. If there are any restore processing failures, view the Tivoli Storage Manager error log for more information.

Tip: The log file is saved to `c:\Program Files\Tivoli\TSM\baclient\dsmerror.log`

Running incremental virtual machine backups using the VStore API

Windows

An incremental backup stores the files that have changed since the last full backup.

Tip: There is no limit to how many full and incremental backups you can take. However, if you do not run a full backup regularly, the size of incremental backups can increase. This scenario ensures that the incremental backups do not get too large.

1. Start the backup of a VM using the following steps:
 - a. Open the backup-archive command line.
 - b. At the command line, enter the following command:

```
backup vm myVMname -vmbackuptype=fullvm -mode=Incremental
```

where *myVMname* is the name of the virtual machine you are backing up. The backup process starts and displays the progress of the backup. The backup is complete when the following results are displayed:

```
Backup VM command complete
Total number of virtual machines backed up successfully: 1
  virtual machine myVMname backed up to nodename NODE
Total number of virtual machines failed: 0
Total number of virtual machines processed: 1
```

2. Verify that you can restore the virtual machines files using the following steps:
 - a. At the command-line interface of the off-host backup server, enter the command:

```
dsmc restore vm my_vm_name -RESTORED
```
 - b. If there are any restore processing failures, view the Tivoli Storage Manager error log for more information.

Tip: The log file is saved to `c:\Program Files\Tivoli\TSM\baclient\dsmerror.log`

Restoring full VM backups

Windows

You can restore a full VMware backup to recreate all of the files of a VMware guest.

You can restore the backup files directly to the VMware server. The restore procedure is different than using the VMware Consolidated Method tools that required you to restore the files to the off-host backup server and use the VMware converter tool before you could restore it to the VMware server.

1. If the full-VM that you are restoring will replace the existing VMware guest, delete the existing full-VM guest before you restore.

Tip: If you restore the full VM backup to a new VMware guest, you do not need to delete the original.

2. Query the Virtual Machine for full VMware backups using the following steps:
 - a. From the off-host backup server, enter the following command:

```
dsmc q vm *
```

The command lists the backups like the following screen:

#	Backup Date	Mgmt Class	Type	A/I	Virtual Machine
1	12/03/2009 03:05:03	DEFAULT	VMFULL	A	vm_guest1
2	09/02/2010 10:45:09	DEFAULT	VMFULL	A	vm_guest11
3	09/02/2010 09:34:40	DEFAULT	VMFULL	A	vm_guest12
4	09/02/2010 10:10:10	DEFAULT	VMFULL	A	vm_guest13
5	12/04/2009 20:39:35	DEFAULT	VMFULL	A	vm_guest14
6	09/02/2010 11:15:18	DEFAULT	VMFULL	A	vm_guest15
7	09/02/2010 02:52:44	DEFAULT	VMFULL	A	vm_guest16
8	08/05/2010 04:28:03	DEFAULT	VMFULL	A	vm_guest17
9	08/05/2010 05:20:27	DEFAULT	VMFULL	A	vm_guest18
10	08/12/2010 04:06:13	DEFAULT	VMFULL	A	vm_guest19
11	09/02/2010 00:47:01	DEFAULT	VMFULL	A	vm_guest7
12	09/02/2010 01:59:02	DEFAULT	VMFULL	A	vm_guest8
13	09/02/2010 05:20:42	DEFAULT	VMFULL	A	vm_guest9

ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

- b. View the output to identify the VMware that you are restoring.
3. Restore the full VMware backup by entering the following command:


```
dsmc RESTORE VM -vmname=my_vmname -datacenter=my_datacenter -host=my_hn -datastore=myPath
```

 See the following parameters to see what you need to substitute for your environment.

my_vmname

This is the display name of the virtual machine

my_datacenter

This is the name of the VMware data center that is defined to the vSphere vCenter.

my_hn

This is the ESX host server name that is defined to the vCenter data center.

myPath

This is the full path location and file name of the volume data and configuration files for the VM backup that you are restoring. You can enter SAN, local storage NAS or iSCSI formatted file paths as defined in vSphere vCenter Data center.

4. When the restore is complete, the virtual machine is powered off. Start the virtual machine from the VMware vCenter to use it.

Backup VM

Windows

Use the **backup vm** command to run a full backup of a specified virtual machine.

The **backup vm** command is used to back up VMware virtual machines from the off-host backup of VMware virtual machine data proxy system..

VMware backup

To specify the name of the virtual machine or to list the virtual machine names for a VMware backup, use the `domain.vmfull` option. If no `vmname` is specified on the command, use the `domain.vmfile` or `domain.vmfull` options.

You can run a full VM backup, that stores a backup copy of complete virtual disk images and configuration information of a virtual machine. Full VM backup enables a complete restore of the complete virtual machine.

You can also specify the `-mode` option to run an incremental or full backup when you use the full VM backup option. File level VM backup provides individual file restore within the virtual machine and incremental backup, although it does not have an easy full machine recovery procedure.

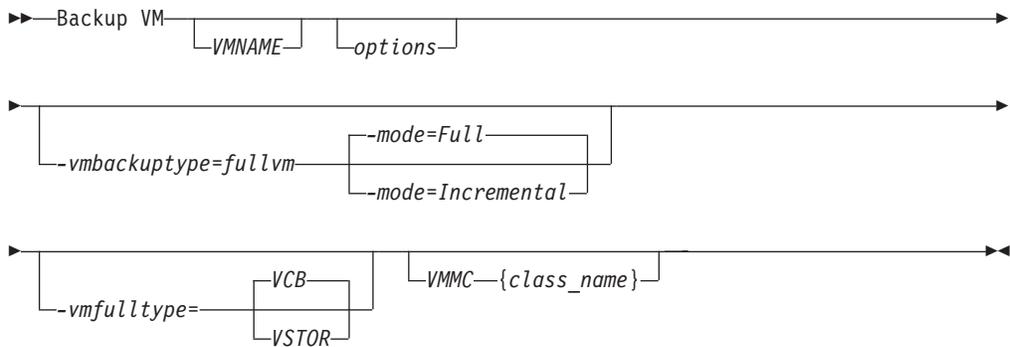
You might want to use a combination of file level VM backups with periodic full VM backups for optimal results.

Tip: You can only use the compression option with the VSTORE full vm backup if the backup is being saved to a storage pool that has been enabled for deduplication.

Supported Clients

This command is valid for Windows that are configured as an off-host backup of VMware virtual machine data proxy.

Syntax



Parameters

Table 15. Backup VM command: Related options

Option	Where to use
<code>vmname</code>	Command line
<code>domain.vmfile</code>	Command line or <code>dsm.opt</code>
<code>domain.vmfull</code>	Command line or <code>dsm.opt</code>
<code>mode</code>	Command line or <code>dsm.opt</code>
<code>vmbackdir</code>	Command line or <code>dsm.opt</code> This option is only supported if you set <code>-vmfulltype=VCB</code>

Table 15. Backup VM command: Related options (continued)

Option	Where to use
vmbacknodelete	Command line or dsm.opt This option is only supported if you set -vmfulltype=VCB
vmbackuptype	Command line or dsm.opt
vmfulltype	Command line or dsm.opt. This option only applies if you specify the option vmbackuptype=fullvm.
VMMC	Command line or dsm.opt. This option only supported if you specify the option vmbackuptype=fullvm or vmfulltype=vstor.
vmbackvcbtransport	Command line or dsm.opt
vmchost	Command line or dsm.opt
vmcpw	Command line or dsm.opt
vmcuser	Command line or dsm.opt

VMware examples

VStore API example commands:

```
dsmc backup vm vm1 -vmbackuptype=file
```

```
dsmc backup vm vm3,vm4 -vmbackuptype=fullvm
```

```
dsmc backup vm vmlocal -vmbackuptype=fullvm
```

To run a file-level virtual machine backup of vm1.example.com using the VMware VirtualCenter machine virtctr.example.com, see the following example:

```
dsmc backup vm vm1 -vmbackuptype=file -vmchost=virtctr
```

Restoring virtual machine data

With Data Protection for VMware, you can perform file-level restore and instant restore of virtual machine data.

- “Mounting snapshots”
- “Tape configuration guidelines” on page 81
- “Restoring files and instant restore of volumes” on page 85
 - “Restoring files (Windows)” on page 87
 - “Using instant restore (Windows)” on page 89
 - “File level restore and instant restore (Linux)” on page 91

Mounting snapshots

You can use Data Protection for VMware Recovery Agent to mount a snapshot and use the snapshot to complete data recovery.

Data Protection for VMware Recovery Agent must be installed and operated from a system that is connected to the Tivoli Storage Manager server through a LAN. Data Protection for VMware does not support operations in a LAN-free environment.

You can use Data Protection for VMware Recovery Agent from either its graphical user interface or from the command line by using the “**mount**” on page 101 command.

Important:

- Mounting a snapshot from the same tape storage pool by two instances of Mount causes one of these results:
 - The second Mount instance is blocked until the first instance is complete.
 - The second Mount instance might interrupt the activity of the first instance. For example, it could interrupt a file copy process on the first Mount instance.
- Avoid concurrent mount sessions on the same tape volume.
- Mount cannot connect to multiple servers or nodes simultaneously.
 - The same snapshot can be mounted more than once.
 - Mount supports a maximum of 20 iSCSI sessions.

When viewing the operating system version of the virtual machine, the Data Protection for VMware Recovery Agent shows the version that was specified when the virtual machine was originally created. As a result, Data Protection for VMware Recovery Agent might not reflect the current operating system.

When a network failure interrupts a mount operation, the volume becomes unstable. A message is issued to the event log. When the network connection is reestablished, another message is issued to the event log. These messages are not issued to the Data Protection for VMware Recovery Agent GUI.

You cannot perform an instant restore or a mount operation for any file system or disk on a virtual machine that is being backed up. You must either wait for the backup to complete, or you must cancel the backup before running an instant restore or a mount operation. These operations are not allowed because the locking mechanism is for a full machine.

In order to prevent possible mount errors, disable the FS Automount option for the Tivoli Storage Manager disks. This option is set in the configuration file used for the HAL daemon. If the FS Automount option is not disabled, the dismount operation might fail.

Data Protection for VMware Recovery Agent on Windows

Data Protection for VMware Recovery Agent can be installed and operated from the following Windows operating systems:

Windows 2003 Service Pack 1 or later

- Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture.
- Supports 32-bit and 64-bit processors.

Windows 2008 Service Pack 1 or later

- Supports the x86 (32-bit), x64 (AMD64 and EM64T) instruction set architecture.
- Supports 32-bit and 64-bit processors.

Windows 2008 R2 or later

- Supports the x64 (AMD64 and EM64T) instruction set architecture.
- Supports 64-bit processors.

Windows XP Professional Edition, Service Pack 2 or later

- Supports the x86 (32 bit) instruction set architecture.
- Supports 32-bit processors.

Windows Vista Service Pack 1 or later

- Supports the x86 (32-bit) and x64 (AMD64 and EM64T) instruction set architecture.
- Supports 32-bit and 64-bit processors.

Windows 7

- Supports the x86 (32 bit) and x64 (AMD64 and EM64T) instruction set architecture
- Supports 32-bit and 64-bit processors.

Note: Data Protection for VMware Recovery Agent can use only snapshots that were created by a Tivoli Storage Manager Backup-Archive Client V6.2.3 or later.

For systems that run Windows 7, Windows Vista, or Windows 2008, Data Protection for VMware Recovery Agent can run in the following two modes:

- When no users are logged in, Data Protection for VMware Recovery Agent runs as a service. The Data Protection for VMware Recovery Agent service enables remote connections through the command line.
- When a user is logged in, Data Protection for VMware Recovery Agent continues to run as a service until you start the Data Protection for VMware Recovery Agent application and use the graphical user interface. When you close the Data Protection for VMware Recovery Agent application and graphical user interface, the Data Protection for VMware Recovery Agent service restarts.

To start Data Protection for VMware Recovery Agent from the Windows Start menu, select **Programs > Tivoli Storage Manager > Data Protection for VMware > Data Protection for VMware Recovery Agent**.

You can use only the Data Protection for VMware Recovery Agent application and graphical user interface when running with administrator login credentials. Only one copy of the Data Protection for VMware Recovery Agent application can be active at any time.

Snapshots can be mounted in either read-only or read/write mode. In read/write mode Data Protection for VMware Recovery Agent saves changes to data in RAM.

Note: If the service is restarted, the changes are lost.

When mounted volumes exist and you start Mount from the Start menu on Windows 7, Windows Vista, or Windows 2008, this message is displayed:

Some snapshots are currently mounted. If you choose to continue, these snapshots will be dismounted. Note that if a mounted volume is currently being used by an application, the application may become unstable. Continue?

When **Yes** is clicked, the mounted volumes are dismounted, even when they are in use.

Setting options

Before proceeding with any operations, click **Settings** in the Data Protection for VMware Recovery Agent GUI to specify the following options:

Virtual Volume write cache

Data changes created during Linux instant restore and mount are saved on a virtual volume in the write cache. By default, the write cache is enabled and specifies the C:\ProgramData\Tivoli\TSM\TDPVMware\mount\ path and

the maximum cache size is 90% of the available space for the selected folder. To prevent the system volume size from becoming full, change the write cache to a path located on a volume other than the system volume.

Folder for temporary files

Specify the path where data changes are saved. The write cache must be located on a local drive and cannot be set to a path on a shared folder. If the write cache is disabled or full, attempting to start an instant restore or mount session on Linux fails.

Cache size

Specify the size of the write cache. The maximum allowed cache size is 90% of the available space for the selected folder.

Data Access

Specify the type of data to be accessed. If you are using an offline device (such as tape or virtual tape library), you must specify the applicable data type.

Storage type

Specify one of the following storage devices from which to mount the snapshot:

Disk/File

The snapshot is mounted from a disk or file. This device is the default.

Tape The snapshot is mounted from a tape storage pool. When this option is selected, it is not possible to mount multiple snapshots or run an instant restore operation.

VTL The snapshot is mounted from an offline virtual tape library. Concurrent mount sessions on the same virtual tape library are supported.

Note: When the storage type is changed, you must restart the service for the changes to take effect.

Read Ahead size (in blocks)

Specify the number of extra data blocks retrieved from the storage device following a read request to a single block. The default values are as follows:

- Disk or file: 64
- Tape or VTL: 1024

The maximum value for any device is 1024.

Read Ahead cache size (in blocks)

Specify the size of the cache where the extra data blocks are stored. The default values are as follows:

- Disk or file: 1000
- Tape or VTL: 75000

Since each snapshot has its own cache, make sure to plan how many snapshots are mounted or restored simultaneously. The cumulative cache size cannot exceed 75000 blocks.

Data Protection for VMware Recovery Agent on Linux

Data Protection for VMware Recovery Agent can be installed and operated from any Red Hat Enterprise Linux 5.2, 5.3, 5.4, 5.5, 5.6, 6.x Server or SuSE Linux Enterprise Server 10, 11 system. Instructions for using Data Protection for VMware Recovery Agent on Linux are available at “File level restore and instant restore (Linux)” on page 91

Note: If a Linux mount operation returns a general error, go to the shell or to the mount log to determine the specific cause of the problem.

If an attempt to remove the connection to the Tivoli Storage Manager server fails, complete the following task:

1. Make sure that there are no mount sessions and no instant restore sessions defined for the Tivoli Storage Manager server. This applies to both the Data Protection for VMware Recovery Agent Windows mount proxy and the Data Protection for VMware Recovery Agent on Linux.
2. If no sessions exist, remove the MountDefinedRepositories.cfg file from the /opt/tivoli/tsm/tdpvmware/mount directory. Restart the Data Protection for VMware Recovery Agent.
3. If sessions exist, contact your IBM support representative.

Related reference

“Supported operating systems for Data Protection for VMware Recovery Agent” on page 7

Tape configuration guidelines

Review these guidelines before attempting backup or restore operations to tape storage.

Preparing for backup

Before attempting a backup to tape, these parameters must be set on the Tivoli Storage Manager server for your tape backups:

1. Define the management class :

```
define mgmtclass <domain name> <policy set name> <mgmtclass name>
```

For example:

```
define mgmtclass tape tape DISK
```

2. Define the copy group:

```
define copygroup <domain name> <policy set name> <mgmtclass name>  
destination=<stgpool name>
```

For example:

```
define copygroup tape tape DISK destination=Diskpool
```

3. Activate the policy set:

```
activate policysset <domain name> <policy set name>
```

For example:

When configuring backup to physical tape, there are additional configuration requirements. You must always keep Tivoli Storage Manager metadata (control files) on disk and the actual virtual machine backup data on tape.

- Use the VMMC option to store the VMware backups (and VMware control files) with a management class other than the default management class.
- Use the VMCTLMC option to specify the management class to use specifically for VMware control files during VMware backups. The management class that you specify overrides the default management class. It also overrides the management class specified by the VMMC option. The VMCTLMC management class must specify a disk storage pool, with no migration to tape. When using Tivoli Storage Manager Backup-Archive Client version 6.2.x, VMCTLMC cannot be configured with the preference editor and must be manually added to dsm.opt. VMCTLMC can be configured manually or with the preference editor when using the Backup-Archive Client version 6.3.
- The VMMC option is always used to control the retention on VM backups. This option applies to both disk and tape configurations. VMCTLMC is not used for the retention of the control files. The control and data files are part of the same grouping and are expired together based on the retention policy of the VMMC option. When both options are set, VMMC is used for data files and VMCTLMC is used for control files.

If a Tivoli Storage Manager server environment uses disk to tape migration, consider the following recommendations:

- Set the disk storage pool MIGDELAY to a value that supports most mount requests to be satisfied from disk. Typical usage patterns indicate that a high percentage of individual file recoveries occurs within few days. For example, usually 3 - 5 days from the time a file was last modified. Therefore, consider keeping data on disk for this brief period to optimize recovery operations.

In addition, if client side deduplication is being used with the disk storage pool, set the MIGDELAY option with a value that considers frequency of full virtual machine backups. The recommendation is not to migrate data from the deduplicated storage pool to tape until at least two full backups can be performed for a virtual machine. When data is moved to tape, it is no longer deduplicated. For example, if full backups are performed weekly, consider setting MIGDELAY to a value of at least 10 days. This setting ensures that each full backup identifies and uses duplicate data from the previous backup before being moved to tape.

- Use a device class file storage pool rather than a DISK device class storage pool. A typical value for a volume size, specified by a device class MAXCAPACITY parameter, would be 8 GB to 16 GB. For the associated storage pool, consider using collocation by file space. Each virtual machine that is backed up is represented as a separate file space in the Tivoli Storage Manager server. Collocating by file space saves the data from multiple incremental backups for a given virtual machine in the same volume (disk file). When migration to tape occurs, collocation by file space locates multiple incremental backups for a given virtual machine together on a physical tape.

Use the **Settings** dialog to set the Tape Mode value.

A backup operation becomes interrupted when a mount or instant restore operation requires the same tape storage simultaneously in use by the backup operation.

Preparing for restore

File restore from physical tape

File restore from a mounted snapshot volume on tape is supported with the following limitations:

- To mount a snapshot volume on physical tape, Data Protection for VMware Recovery Agent mount must be operating in Tape Mode. In this mode, only one virtual machine snapshot volume can be mounted at a time from a Data Protection for VMware Recovery Agent. The limit is one disk when mounting an iSCSI target or one volume when creating a virtual volume from a selected partition. You can have multiple Tivoli Storage Manager tape storage pool volume mounts for different virtual machine snapshots. Create this scenario by installing the Data Protection for VMware Recovery Agent on multiple physical machines or on multiple virtual machine guests. Enable Tape Mode by selecting Storage Type=Tape in the Data Protection for VMware Recovery Agent GUI Settings menu.
- While Data Protection for VMware Recovery Agent mount does not prevent an attempt to mount a snapshot located on physical tape, performance can vary significantly and might be severely affected. Data Protection for VMware Recovery Agent mount does not control the way data is accessed on tape. Data access patterns can be random and cause extensive challenges associated with tape positioning operations. For Linux virtual machines, Tape Mode is only supported when directly using the iSCSI initiator. Tape Mode is not supported from the Linux Data Protection for VMware Recovery Agent user interface.

While a limited number of files can be recovered from a mounted volume snapshot located on physical tape, consider performing a full virtual machine restore from physical tape. If you must recover a large quantity of data or many files.

Typically the performance associated with a file restore from a VTL is quicker than physical tape. But, performance delays can be encountered. For example, mount might require a virtual tape volume that is in use by another restore operation. Tape volumes in a VTL cannot be shared. In this case, mount processing is delayed until the restore operation using the virtual tape volume completes.

When a mount or instant restore operation requires the same tape storage simultaneously in use by a backup operation, the backup operation becomes interrupted.

Instant restore from physical tape

The Data Protection for VMware Recovery Agent instant restore operation is not supported for snapshot volumes located on physical tape.

The failure behavior for instant restore depends on the mode of operation:

- Tape mode: Instant restore fails when the restore operation is selected.
- Non-tape mode: Instant restore fails when an attempt is made to access data on a tape volume.

This failure can occur when the user interface tries to show partition information or when the operation is restoring data. The latter condition

occurs only if disk to tape migration is being used, and part of the snapshot data (the partition information) is located on disk and some of the actual snapshot data is located on tape.

If you try to run instant restore in Tape Mode, the following message is shown:

Instant Restore is not supported in Tape Mode.

If you try to run instant restore while not in Tape Mode, and the data is located on tape, the following message is shown:

An error occurred while reading snapshot data from server. See log for details.

Configuring systems for iSCSI mount

This procedure describes how to configure systems used during an iSCSI mount operation. The snapshot is mounted from Tivoli Storage Manager server storage.

Be aware of these requirements before proceeding with this task:

- Make sure Microsoft iSCSI Initiator is installed.
- If a volume spans several disks, you must mount all the required disks. When mirrored volumes are used, mount only one of the mirrored disks. Mounting one disk prevents a time-consuming synchronization operation.
- If multiple dynamic disks were used on the backup system, these disks are assigned to the same group. As a result, Windows Disk Manager might consider some disks as missing and issue an error message when you mount only one disk. You can ignore the message as the data on the backed up disk is still accessible, unless some of the data is located on the other disk. This issue can be solved by mounting all the dynamic disks.

Complete these tasks to configure the systems used during an iSCSI mount operation:

1. Open port 3260 in the LAN firewall and the Windows client firewall. Record the iSCSI Initiator name as shown in the iSCSI Initiator configuration window of the Control Panel. For example:
`iqn.1991-05.com.microsoft:hostname`
2. Complete these tasks on the system running the Data Protection for VMware Recovery Agent (or iSCSI target):
 - a. Start the Data Protection for VMware Recovery Agent. Complete the Select TSM server and Select snapshot dialogs and click **Mount**.
 - b. In the Choose mount destination dialog, select Mount an iSCSI target.
 - c. Create a target name. Make sure that it is unique and that you can identify it from the system that runs the iSCSI initiator. For example:
`iscsi-mount-tsm4ve`
 - d. Enter the iSCSI Initiator name recorded in Step 1 and click **OK**.
3. Verify that the volume you just mounted is displayed in the Mounted Volumes field.
4. Locate and start the iSCSI Initiator program on the initiator system selected in Step 1:
 - On Windows Server 2003:

- a. In the Discovery tab, click **Add** in the Target Portals dialog. Enter the TCP/IP address of the Data Protection for VMware Recovery Agent (iSCSI target) used in Step 2. Click **OK**.
- b. Verify that the iSCSI target is visible in the Targets tab. Click **Refresh** if it is not visible.
- c. Select the target and click **Log On** to connect the iSCSI virtual volume.
- On Windows 7 and Windows Server 2008:
 - a. In the Targets tab, enter the TCP/IP address of the Data Protection for VMware Recovery Agent (iSCSI target) used in Step 2 in the Target: dialog. Click **Quick Connect**.
 - b. The Quick Connect dialog shows a target that matches the target name specified in Step 2c. If it is not already connected, select this target and click **Connect**.
5. On the initiator system, go to **Control Panel-> > Administrative Tools-> > Computer Management-> > Storage > Disk Management**.
 - a. If the mounted iSCSI target is listed as Type=Foreign, right-click the Foreign Disk and select Import Foreign Disks. The Foreign Disk Group is selected. Click **OK**.
 - b. The next screen shows the type, condition, and size of the Foreign Disk. Click **OK** and wait for the disk to be imported.
 - c. When the disk import completes, press **F5** (refresh). The mounted iSCSI snapshot is visible and contains an assigned drive letter. If drive letters are not automatically assigned, right-click the required partition and select Change Drive Letters or Paths. Click **Add** and select a drive letter.
6. Open Windows Explorer (or other utility) and browse the mounted snapshot for a file-level recovery operation. See “File level restore and instant restore (Linux)” on page 91 for instructions.

Restoring files and instant restore of volumes

With the stored snapshots, you can recover data that is backed up. You can restore files and perform the instant restore of volumes.

Note: Do not attempt to change a Tivoli Storage Manager node password while running a file level restore or an instant restore from snapshots stored in that node. The results are unpredictable.

Restoring files

Administrators can use Data Protection for VMware Recovery Agent for efficient file level restores and to minimize downtime by mounting snapshots to virtual volumes.

The virtual volume can be viewed by using any file manager, for example Windows Explorer. The directories and files in the snapshot can be viewed and managed like any other file. If you edit the files and save your changes, after you unmount the volume, your changes are lost because the changed data is held in memory and never saved to disk. Because the changes are written to memory, Data Protection for VMware Recovery Agent can use a large amount of RAM when working in read/write mode.

You can copy the changed files to another volume before performing an unmount.

You can select *read only* as a mounting option.

Data Protection for VMware Recovery Agent mounts snapshots from the Tivoli Storage Manager server.

Data Protection for VMware Recovery Agent can be used for the following tasks:

- Recovering lost or damaged files from a backup
- Mounting a virtual machine guest volume and creating an archive of the virtual machine guest files
- Mounting database applications for batch reports

From the Data Protection for VMware user interface, you must first dismount any snapshots before you click **Remove**. The remove operation fails if there are active mount and restore sessions in the Linux or Windows Mount machines. You cannot remove the connection to server when you are performing a file restore or an instant restore from that server. You must first dismount all virtual devices and stop all instant restore sessions before you disconnect from a server. If you do not do so, the connection is not removed.

You must unmount all virtual volumes before uninstalling Data Protection for VMware Recovery Agent. Otherwise, these mounted virtual volumes cannot be unmounted after Data Protection for VMware Recovery Agent is reinstalled.

To use Data Protection for VMware Recovery Agent for file-level recovery of data that is stored on tape, it is recommended that the data be moved to disk or file storage. From Tivoli Storage Manager, you can use the **QUERY OCCUPANCY** command to see where the data is stored. You can then use the **MOVE NODEDATA** command to move this data back to disk or file storage.

For more information about these commands, see the Tivoli Storage Manager Information Center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>

Copying files from a mounted snapshot tape storage pool performs more slowly than does a snapshot that is on disk. It might take less time to move the backup data for a virtual machine from tape to disk before performing a file restore. In such cases, you would use the **MOVE NODATA** command with the **FILESPACE** option. This approach might be better for including file restore if there are many files on a badly fragmented volume.

Note: Mounting a snapshot from the same tape storage pool by two instances of Mount can cause one of these results:

- The second Mount instance is blocked until the first instance is complete.
- Both mounts succeed, but the performance is poor.

When restoring data from a mirrored volume, mount only one of the disks that contains the mirrored volume. Mounting both disks causes Windows to attempt a resynchronization of the disks. However, both disks contain a different timestamp if mounted. As a result, all data is copied from one disk to the other disk. This amount of data cannot be accommodated by the virtual volume. When you must recover data from a volume that spans two disks, and those disks contain a mirrored volume, complete these steps:

1. Mount the two disks.
2. Use the iSCSI initiator to connect to the first disk.
3. Use Windows Disk Manager to import this disk. Ignore any message regarding synchronization.
4. Delete the mirrored partition from the first (or imported) disk.

5. Use the iSCSI initiator to connect to the second disk.
6. Use Windows Disk Manager to import the second disk.

Both volumes are now available.

Instant restore

You can use instant restore to use data on a volume that is being restored, while the restore operation is in progress. For this reason, less downtime is required before a recovered volume can be used.

Instant restore works only with local volumes. Local volumes must have an assigned drive letter.

The file system on the destination folder cannot be a FAT file system. If you plan to restore into a FAT volume, you must format it as NTFS before attempting an instant restore.

You can complete an instant restore of a volume in a supported clustered environment. While instant restore process is running, you can access the volume. Other volumes in the cluster should not be affected, and you can work with the cluster, and with that volume, in parallel. During the instant restore, the disk being restored cannot fail over if the node fails.

If a system is shut down while instant restore is in progress, the instant restore automatically continues from the same point when power is restored. For Linux, the instant restore progress can be slow in some situations. This can be due to such things as connectivity issues, very large volumes, restore from tape, or high I/O operations. The recovery process identifies no progress was made on the MDADM device (`/dev/md0`) for the last 5 minutes and responds to it as a potential problem. As a result, a cleanup is performed, and the restore restarts itself, continuing from the point at which it stopped. The cleanup can cause I/O operations to fail. To fix this problem, you can increase the recovery process period by editing the `/etc/crn.d/tsmmount_recover` file. Increase the recovery process period to 10 minutes or to 20 minutes if the volume is very large, or the I/O is intense.

Instant restore destination volumes must be either on basic disks, or simple volumes on dynamic disks. Destination volumes cannot be spanned volumes, mirrored volumes, or RAID-5 volumes. You can use a basic disk as a destination volume and then convert the basic disk to a dynamic disk.

Restoring files (Windows)

You can use Data Protection for VMware Recovery Agent for efficient file level recovery and to minimize downtime by mounting snapshots to virtual volumes. On supported Windows operating systems, file-level recovery is supported from snapshots of NTFS, FAT, or FAT32 volumes.

The mount function cannot be used to mount snapshots of dynamic disks as a virtual volume. Only partitions from an MBR-based, basic disk can be mounted as virtual volumes. File-level recovery from GPT, dynamic, or any other non-MBR or non-basic disk is possible by creating a virtual iSCSI target and using an iSCSI initiator to connect it to your system.

| **Important:** The ACL values associated with the folders and files that are restored
| in a file level recovery operation are not transferred to the recovered files. In order
| to maintain ACL values, use the XCOPY command when copying files from the
| target in Step 6.

To run a file level recovery for a Windows system, complete the following steps:

1. Log on to the system where you want to restore files. Data Protection for VMware Recovery Agent must be installed on the system.
2. Start Data Protection for VMware Recovery Agent
 - **For Windows 7, Windows Vista, or Windows 2008 only:** Select **All Programs > Tivoli Storage Manager > Data Protection for VMware > Data Protection for VMware Recovery Agent**
 - **For all other supported Windows systems:** From the Microsoft Windows taskbar area, click the Data Protection for VMware Recovery Agent icon.
3. Connect to a Tivoli Storage Manager server by specifying the server address, port, target node, and password. Although the Select TSM server list appears to contain multiple servers, this list contains a maximum of one server only. The target node is where the backups are located. You can manage the level of access to the target node data by specifying a different node name in the Node access method section:

Asnodename

Select this option to use the asnodename feature to access the target node. Although backups are located in the target node, you can use a proxy node that is granted proxy authority to the target node. Typically, the proxy node is created by the Tivoli Storage Manager administrator. After selecting the check box, enter the name of the *target* node in the Target node field, the name of the *proxy* node in the Authentication Node field, and the password for the proxy node in the Password field. When using this method, only the proxy node password should be known. As a result, the target node password is not exposed.

Fromnode

Select this option to use a node that provides query and restore access only to the target node data. After selecting the check box, enter the name of the *target* node in the Target node field. Then, in the Authentication Node and Password fields, enter the node name and password of a node that was granted access only to the snapshot data of specific virtual machines stored in the target node. You can grant access to the node by using the Tivoli Storage Manager Backup-Archive Client **set access** command. This option is used to limit the number of virtual machines available for restore to a certain group, instead of having all virtual machines from the target node available. If you must revoke the authentication node access to the storage node, issue the **delete access** command.

Note: When using this method, the snapshot data is not protected from expiration on the server. As a result, instant restore is not supported in this method.

Direct Select this option to authenticate directly to the target node.

Data Protection for VMware Recovery Agent queries the specified server for a list of protected virtual machines, and shows the list.

4. Select a virtual machine from the list. Data Protection for VMware Recovery Agent queries the server for a list of snapshots available for the specified virtual machine.

A virtual machine might appear in the virtual machines list, but if you select it, the snapshots list might be empty. This situation occurs because of one of the following reasons:

- No snapshots have yet completed successfully for that virtual machine.
 - The **Fromnode** option was used and the specified node is not authorized to restore the selected virtual machine.
5. Select the required snapshot by selecting the date and disk and click **Mount**.
 6. In the Mount Destination dialog, check **Mount the Following Partition**. A list of partitions available on the selected disk is shown. For each partition, its size, label, and file system type are displayed.
 - If the disk is not MBR-based, an error message is displayed.
 - By default, only partitions that can be used for file-level restore are displayed.
 - To display all partitions that existed on the original disk, clear the **Show only mountable partitions** check box.
 - When multiple dynamic disks exist in the virtual machine guest, all dynamic disks that belong to the same group must be mounted and connected to the iSCSI initiator. Otherwise, Windows Disk Manager might consider some disks as missing and issue an error message. You can ignore the message as the data on the backed up disk is still accessible.

Tip: If the **Mount the Following Partition** checkbox is disabled, you can use the **Mount as an iSCSI target** feature described in Step 9.

7. Select the required partition. Partitions formatted using unsupported file systems cannot be selected.
8. Specify a drive letter or an empty folder as a mount point for the virtual volume.
9. If the device is an iSCSI disk, use the iSCSI initiator to discover and log on to the iSCSI target. You can view and copy files from the target.
10. Select **Mount as an iSCSI device**, and enter the initiator name and the target name. Click **OK**.

Using instant restore (Windows)

With instant restore, you can restore a volume and almost immediately use the restored volume. Less downtime is required before a recovered volume can be used because you can use data on the disk while the restore is in progress.

Instant restore is available only from Data Protection for VMware snapshots of disks of the MBR type. The volume format of volumes on those disks must be NTFS, FAT, or FAT32. However, instant restore to a destination partition on FAT volumes is not supported. As a result, if you plan to restore to a destination partition that is formatted as FAT, you must format the partition as NTFS before attempting a restore. In addition, when selecting a destination volume for instant restore, make sure that the volume resides on a physical disk, and not on a virtual iSCSI disk.

- Restoring a volume involves overwriting data on the existing storage volume. After the restore begins, the current volume contents are permanently erased. Before you start the restore, verify that the correct volume is selected, and that there are no open handles or processes using that volume.

- The restore operation fails if there are open files or applications that are running on the target restore volume. Selecting **Ignore open handles on the destination volume** causes Data Protection for VMware to ignore the open files and applications that are running on the destination volume. This situation can cause a problem with applications and loss of data in files that are open on the target volume.

To perform an instant restore, complete the following steps:

1. Start Data Protection for VMware Recovery Agent.
 - **For Windows 7, Windows Vista, and Windows 2008 only:** Select **All Programs > Tivoli Storage Manager > Data Protection for VMware > Data Protection for VMware Recovery Agent**
 - **For all other supported Windows systems:** From the Microsoft Windows taskbar area, click the Data Protection for VMware Recovery Agent icon.
2. In the Data Protection for VMware Recovery Agent window, select the Tivoli Storage Manager server to use as the source. Although the Select TSM server list appears to contain multiple servers, this list contains a maximum of one server only. Data Protection for VMware Recovery Agent queries the Tivoli Storage Manager server for a list of protected virtual machines and displays the list.
3. You can manage the level of access to the target node data by specifying a different node name in the Advanced options dialog:

Asnodename

Select this option to use the asnodename feature to access the target node. Although backups are located in the target node, you can use a proxy node that is granted proxy authority to the target node. Typically, the proxy node is created by the Tivoli Storage Manager administrator. After selecting the check box, enter the name of the *target* node in the Target node field, the name of the *proxy* node in the Authentication Node field, and the password for the proxy node in the Password field. When using this method, only the proxy node password should be known. As a result, the target node password is not exposed.

Fromnode

Select this option to use a node that provides query and restore access only to the target node data. After selecting the check box, enter the name of the *target* node in the Target node field. Then, in the Authentication Node and Password fields, enter the node name and password of a node that was granted access only to the snapshot data of specific virtual machines stored in the target node. You can grant access to the node by using the Tivoli Storage Manager Backup-Archive Client **set access** command. This option is used to limit the number of virtual machines available for restore to a certain group, instead of having all virtual machines from the target node available. If you must revoke the authentication node access to the storage node, issue the **delete access** command.

Note: When using this method, the snapshot data is not protected from expiration on the server. As a result, instant restore is not supported in this method.

Direct Select this option to authenticate directly to the target node.

4. Select a virtual machine, date, time, and disk and click **Restore**.
5. Data Protection for VMware Recovery Agent displays a list of partitions available on the selected disk. For each partition, its size, label, and file system

type are displayed. Select the required partition. By default, only partitions that can be restored are displayed. To display all the partitions that are available on one or more disks, clear the **Show only restorable partitions** check box. Select the required partition from the list.

Note:

- Drive letters are not displayed.
 - If a disk cannot be parsed, an error message is displayed and the **Instant Restore** dialog is closed. For example, this occurs when the disk is dynamic or a GUID partition table (GPT).
6. Select the destination partition into which the data is to be restored. The destination location size must be equal or larger than the source size.
 7. Click **Restore**.
 8. A confirmation message is displayed. Verify the information and click **Yes**. The restore process begins. In the instant restore section, you can see the status of the restore process. When the status changes to restoring, the volume is available for use.

Use the **Max CPU** slider to adjust the processor usage for the restore process.

To cancel the restore process, select the instant restore session that is in progress and click **Abort**. All data on the target drive is lost. You can click **Abort All** to cancel all processes. If you stop an instant restore without clicking **Abort** or **Abort all**, the restored volume is displayed as a valid volume, but the data on the volume is invalid. The data is invalid because the data was partially restored, but the restore process did not have time to complete, and the shutdown was abnormal.

If the service is stopped while instant restore is running, the volume appears to be a valid volume. Trying to access the area of the volume that is not yet restored fails, and the data appears corrupted. After the service restarts, the restore process continues, and the data appears valid. If a power failure occurs during instant restore, after the machine boots up, the volume appears to be unformatted. After the service starts, the instant restore process resumes, and the volume appears valid.

A temporary problem might prevent the session from running. For example, a network problem might cause a temporary loss of access to the Tivoli Storage Manager server. In that case, the instant restore session pauses. To continue to the restore process after the pause, select the appropriate line in the instant restore list and click **Resume**. During the period when the session is paused, the parts of the volume that are not yet restored are inaccessible.

You can use instant restore to restore a simple volume that is located on a dynamic disk. The destination volume can be a dynamic disk; however, the source volume cannot be a dynamic disk. This restore might cause the disk status to change to *Online (Errors)*. In addition, the status of all volumes on the disk might change to *At Risk*. This change in disk status can occur when network traffic is too heavy for instant restore to operate. In this situation, the volumes are online and mounted. You can return the disk and volume status to normal by going to the Computer Management Console. Right-click the disk; then, click **Reactivate Disk**.

File level restore and instant restore (Linux)

Data Protection for VMware Recovery Agent on Linux is used to restore individual files (file level restore) or volumes (instant restore). Unlike a conventional volume

restore, instant restore provides access to volume contents while the restore process is in progress. Less downtime is required before a recovered volume can be used. After you start an instant restore, you can use data on the disk while the restore is in progress.

Configuring Data Protection for VMware Recovery Agent for restore operations (Linux):

Data Protection for VMware Recovery Agent requires specific application settings, environment conditions, and configuration tasks be completed before attempting a restore operation.

These environment requirements must exist before using Data Protection for VMware Recovery Agent on Linux:

- The Tivoli Storage Manager command line must be available on a Windows computer.
- Data Protection for VMware Recovery Agent is available on a Windows system. This system must be accessible from the computer where the command line is installed. Alternatively, Data Protection for VMware Recovery Agent and the command line can be installed on the same computer.
- Data Protection for VMware Recovery Agent must be able to access the IBM Tivoli Storage Manager storage pool. Data Protection for VMware Recovery Agent exposes snapshots as iSCSI targets. Therefore, the snapshots must be accessible to the target Linux machine.
- For the iSCSI to work for Linux mount and restore operations, the iSCSI port must be open on any firewall between the machine running Windows mount, the iSCSI target and the machine performing the restore, the iSCSI initiator. The iSCSI default port is 3260.
- Ensure that your environment consists of all prerequisite applications.
- When performing a Linux instant restore or mount, Data Protection for VMware Recovery Agent saves changes to data on a virtual volume in the write cache. The path is `C:\Documents and Settings\All Users\Application Data\Tivoli\TSM\tdpvmware\mount`, and the size is set to a maximum of 90% of the available space. These settings can be configured by clicking settings in the main Data Protection for VMware Recovery Agent window. The write cache must be located on a local drive and cannot be set to a path on a shared folder. You cannot enable or disable the write cache from the UI or from configuration files. It is recommended that you specify the write cache location in a non-system folder on a local disk.
- In order to prevent the recovery process from mounting the device, stop the cron daemon. For example:

```
/etc/init.d/cron stop
```

Start the cron daemon when the processing completes.

- There is a limitation when restoring several volumes to several partitions on the same disk. Only one volume can be active. The other volumes remain in the DELAYED state. Their progress as shown on the user interface remains at 0% until they move out of the DELAYED state after the active synchronization completes.

Restriction: Exposing snapshots as iSCSI targets

When a snapshot of a dynamic disk is exposed to its original machine, the UUIDs become duplicated. This duplication negatively affects how the volume structure is understood. Likewise when a snapshot of a GPT disk is exposed to its original

machine, the GUIDs become duplicated. This duplication also negatively affects how the disk structure is understood. As a result, expose dynamic disks and GPT disks to a machine other than the original machine. For example, expose these disk types to a proxy machine, unless the original disks no longer exist.

This task guides you through configuration steps required to use Data Protection for VMware Recovery Agent.

1. Log on to the Linux system with root user authority. Data Protection for VMware Recovery Agent must be installed on this Linux system.
2. Start Data Protection for VMware Recovery Agent by clicking the Data Protection for VMware Recovery Agent icon on the desktop or running a script from the shell prompt. The first time you access Data Protection for VMware Recovery Agent, the Settings dialog displays. You must enter the following configuration information to proceed:

- **Command line**

- a. Enter the host name or IP address of the computer where the command line is installed.
- b. Enter the login ID that is used for the Secure Shell (SSH) user.

Tip: This login ID is for the Windows system where both the command line and SSH are installed. This system uses SSH to communicate with Data Protection for VMware Recovery Agent on the Linux system. Make sure this login ID uses a host name convention defined in the SSH `known_hosts` file.

- **Data Protection for VMware Recovery Agent** Enter the host name or IP address of the Windows system where Data Protection for VMware Recovery Agent is installed. Click **OK** to save these values and return to the Data Protection for VMware Recovery Agent window.
3. Use the Select TSM server list to identify the server to use as the source. Although this list appears to contain multiple servers, it contains a maximum of one server only. The Tivoli Storage Manager must already be configured and accessible to Data Protection for VMware Recovery Agent.

- Enter the following information:

Server address

Enter the IP address or host name of the Tivoli Storage Manager.

Server port

Enter the port number used for TCP/IP communication with the server. The default port number is 1500.

Asnodename

Select this option to use the asnodename feature to access the target node. Although backups are located in the target node, you can use a proxy node that is granted proxy authority to the target node.

Typically, the proxy node is created by the Tivoli Storage Manager administrator. After selecting the check box, enter the name of the *target* node in the Target node field, the name of the *proxy* node in the Authentication Node field, and the password for the proxy node in the Password field. When using this method, only the proxy node password should be known. As a result, the target node password is not exposed.

Fromnode

Select this option to use a node that provides query and restore access only to the target node data. After selecting the check box,

enter the name of the *target* node in the Target node field. Then, in the Authentication Node and Password fields, enter the node name and password of a node that was granted access only to the snapshot data of specific virtual machines stored in the target node. You can grant access to the node by using the Tivoli Storage Manager Backup-Archive Client **set access** command. This option is used to limit the number of virtual machines available for restore to a certain group, instead of having all virtual machines from the target node available. If you must revoke the authentication node access to the storage node, issue the **delete access** command.

Note: When using this method, the snapshot data is not protected from expiration on the server. As a result, instant restore is not supported in this method.

Direct Select this option to authenticate directly to the target node.

4. Click **OK** to save these values and return to the Data Protection for VMware Recovery Agent window.
5. Click **Refresh** to display the most current data from the Tivoli Storage Manager server.

Data Protection for VMware Recovery Agent is now properly configured and ready for restore operations. Use Data Protection for VMware Recovery Agent to accomplish a file level restore or an instant restore operation.

File-level restore (Linux):

File-level restore on Linux is available from snapshots of disks that use any of the following:

- MBR-style partition tables.
- Partitions formatted using one of these file systems:
 - NTFS
 - FAT
 - FAT32
 - EXT2
 - EXT3
 - EXT4
 - ReiserFS

File restore from snapshots of GPT-based disks, dynamic disks, and LVM partitions is possible by using iMount technology. This technology exposes an iSCSI target. iSCSI initiator commands are then used to discover the iSCSI target.

Be aware of these considerations before attempting a file-level restore on Linux:

- If, during a file restore on a Linux machine, the iSCSI target is dismounted, the file level restore fails, and the restored data is left uncompleted.
- The tasks described in “Configuring Data Protection for VMware Recovery Agent for restore operations (Linux)” on page 92 must be completed before attempting a file level restore.
- This procedure assumes that you are logged on to the Linux system with root user authority and the Data Protection for VMware Recovery Agent GUI is available.

- SUSE Linux Enterprise Server 10 and Enterprise Server 11 require all iSCSI devices to be unmounted before rebooting or shutting down the system.
- When a mount operation of an LVM volume fails, issue the `lvscan` command in order to verify whether the volume is active. If the volume is inactive, issue the `vgchange -a y volume_group` command to make the volume available. Then, retry the mount operation.

Important: The ACL values associated with the folders and files that are restored in a file level restore operation are not transferred to the restored files. In order to maintain ACL values, use the `cp -p` command when copying files.

This task describes how to use Data Protection for VMware Recovery Agent to restore a snapshot volume (file level) on a Linux system.

1. Identify the Tivoli Storage Manager server where the snapshots are stored. Specify the server address, port, node, and password. Although the Select TSM server list appears to contain multiple servers, this list contains a maximum of one server only. Data Protection for VMware Recovery Agent queries the server for a list of protected virtual machines and displays the list.
2. Select a virtual machine from the list. Data Protection for VMware Recovery Agent queries the server for a list of snapshots available for the specified virtual machine. A virtual machine might be displayed in the virtual machines list, but if you select it, the snapshots list might be empty. This situation might occur for one of the following reasons:
 - No snapshots have yet completed successfully for that virtual machine.
 - The node used for authentication was not granted permission to restore the selected virtual machine.

Tip: To quickly locate the required virtual machine from the available virtual machine list, type the first few letters of the virtual machine name.

3. Select the required snapshot by selecting the date and disk. Data Protection for VMware Recovery Agent displays a list of partitions available on the selected disk. For each partition, size, label, and file system type are displayed. By default, only mountable partitions are displayed. To display all partitions, clear the **Show only mountable partitions** check box.

Note: Mount points are not displayed.

4. Select the required partition.
5. Select a path where the virtual volume will be mounted.
6. Click **OK** to start the file level restore.

After the mount process is completed successfully, a new entry is displayed in the **Mounted Volumes** field. For example:

```
/mnt is mount of [tsm-ba-1@tsm-ve-1]-[vm-1]-[2010-Mar-24
10:10:10]-[Hard Disk 1]-[Partition 0]
```

Instant restore (Linux):

Before attempting an instant restore on Linux review the following information:

- Multiple instant restore sessions to different target disks run in parallel. However, multiple instant restore sessions to different target partitions on the same disk do not run in parallel. As a result, the first instant restore session must complete before the next Instant Restore session begins.

- The tasks described in “Configuring Data Protection for VMware Recovery Agent for restore operations (Linux)” on page 92 must be completed before attempting an instant restore.
- This procedure assumes that you are logged on to the Linux system with root user authority and that the Data Protection for VMware Recovery Agent user interface is available.
- SUSE Linux Enterprise Server 10 and Enterprise Server 11 require all iSCSI devices to be unmounted before rebooting or shutting down the system.
- Instant restore to LVM partitions is not supported.
- Instant restore is available for snapshots of disks that use MBR-style partition tables only. The partition used as the source for an instant restore operation must be formatted by using one of the following file systems:
 - NTFS
 - FAT
 - EXT2
 - EXT3
 - EXT4
 - ReiserFS
- When selecting a destination volume for instant restore, make sure that the volume resides on a physical disk, and not on a virtual iSCSI disk.

This task guides you through how to use Data Protection for VMware Recovery Agent to restore a snapshot volume (instant restore) on a Linux system.

1. Identify the Tivoli Storage Manager server where the snapshots are stored. Specify the server address, port, node, and password. Data Protection for VMware Recovery Agent queries the server for a list of protected virtual machines and displays the list.
2. Select a virtual machine from the list. Data Protection for VMware Recovery Agent queries the server for a list of snapshots available for the specified virtual machine.

Tip: To quickly locate the required virtual machine from the available virtual machine list, type the first few letters of the virtual machine name.

3. Select the required snapshot by selecting the date and disk. Data Protection for VMware Recovery Agent displays a list of partitions available on the selected disk. For each partition, its size, label, and file system type is displayed. By default, only restorable partitions are displayed. To display all partitions, clear the **Show only restorable partitions** check box.

Note: Mount points are not displayed.

4. Select the required partition.
5. Select the destination partition into which the data is to be restored, by selecting either a mount point or a block device. If you specify both, ensure that the block device is mounted on the specified mount point.
6. Click **OK**. The restore process starts. After a short initialization period, the volume is available for use while the restore process runs in the background and until the volume is completely restored.

Restoring to the same volume again:

If you plan to restore another snapshot into the same target volume, complete one of the following steps:

- Restart the Linux system.
- Manually stop the mirror device and mount the restored volume.
For example, in the following procedure sdc1 is the target block device and md0 is the mirror device:
 1. Issue the command: `umount /dev/md0`.
 2. Issue the command: `mdadm --stop /dev/md0`.
 3. Issue the command: `mount /dev/sdc1 /restoredVolume`.

Checking the file system

After the instant restore completes, you can verify the file system restored volume by using the `fsck` file system utility:

1. Unmount the RAID device by issuing this command: `umount /dev/md0`
2. Type in the `fsck` command to run the file system check.

Responding to a timeout during a file level restore or an instant restore (Linux):

During a file level restore or an instant restore, a timeout might occur. If a timeout does occur, the user interface displays a message saying that manual intervention might be needed.

Follow this procedure if the timeout occurs during a file level restore (Mount operation):

1. Stop the recovery process by commenting out the one line in `/etc/cron.d/tsmmount_recover` Ensure that the recovery process is not running by issuing this command:
`ps -ef | grep tsmRecover`
2. Ensure that the required snapshot is mounted in the Data Protection for VMware Recovery Agent Windows backup server. If the snapshot is not mounted, mount it manually.
3. Ensure that the iSCSI target is connected to the Linux machine by using **iscsiadm -m session**. If the iSCSI target is not connected to the Linux machine, perform a manual login to the target by issuing this command:
`iscsiadm -m discovery -t sendtargets -p <windows_server_ip> --login`
4. Mount the iSCSI device locally by using the Linux **mount** command. For example:
`mount /dev/sde1 /Mount1`

Ensure that you use the same mount point as was requested in the user interface. This operation can take long time to complete depending on the consistency of the snapshot.

5. When the mount operation is completed, uncomment the line in `/etc/cron.d/tsmmount_recover`.
6. When you want to unmount, do so first locally by using **umount**. Then use the Linux user interface to unmount the session.

For an instant restore operation the manual intervention is to halt and then retry the restore session.

Chapter 8. Recovery Agent command-line interface

The Recovery Agent CLI lets you access most Data Protection for VMware functions. The Recovery Agent CLI can be viewed as a command-line API to the Data Protection for VMware Recovery Agent. Changes completed with the Recovery Agent CLI to the Data Protection for VMware Recovery Agent take effect immediately.

You can use the Recovery Agent CLI to manage only one system running the Data Protection for VMware Recovery Agent.

Starting the Recovery Agent command-line interface

Before you can start and use the Recovery Agent CLI from a supported Linux operating system, you need to complete the software prerequisites detailed in “Software requirements and prerequisites” on page 13.

To start the Recovery Agent CLI, complete the following steps:

1. From the Windows Start menu, click **Programs > Tivoli Storage Manager > Data Protection for VMware > Data Protection for VMware Recovery Agent**.

2. In the command prompt window, enter one of the following commands:

- To run the Recovery Agent CLI:

```
TDPVMwareShell.exe -c command type tag parameter
```

- **Windows** To display the help for the Recovery Agent CLI:

```
TDPVMwareShell.exe -h
```

- **Linux** To display the help for the Recovery Agent CLI:

```
TDPVMwareShell.exe -h dump
```

For example, this command displays detailed help for the mount Recovery Agent CLI: **Linux**

```
TDPVMwareShell.exe -h mount dump
```

The Recovery Agent CLI on Linux requires a proxy node for the Windows system where the Data Protection for VMware Recovery Agent is running. All Recovery Agent CLI commands should be issued from this Windows system.

Recovery Agent command-line interface overview

When you use the commands, some parameters are not required. See the following sections for details regarding required parameters.

For the parameters that are not required and not entered, default values are used. Parameters with spaces must be enclosed in quotation marks. For example, if you want to use the *Accounting, Daily* parameter, type "Accounting, Daily".

To read a syntax diagram for entering a command, follow the path of the line. Read from left to right, and from top to bottom, and use the following guidelines:

- The >>- character sequence indicates the beginning of a syntax diagram.

- The --> character sequence at the end of a line indicates that the syntax diagram continues on the next line.
- The >-- character sequence at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The -->< character sequence indicates the end of a syntax diagram.

Symbols

Enter these symbols exactly as they are displayed in the syntax diagram:

*	Asterisk
{ }	Braces
:	Colon
,	Comma
=	Equal sign
-	Hyphen
()	Parentheses
.	Period
	Space
"	Quotation mark
'	Single quotation mark

Variables

Italicized lowercase items such as *<variable_name>* indicate variables. In this example, you can specify a *<variable_name>* when you enter the **cmd_name** command.

▶▶ --cmd_name--<variable_name>————▶▶

Required choices

When two or more items are in a stack and one of them is on the line, you must specify one item. In the following example, you must choose either *A*, *B*, or *C*:

▶▶ --cmd_name--
A
B
C
————▶▶

Optional choices

When an item is below the line, that item is optional. In the following example, you can select either *A* or nothing at all:

▶▶ --cmd_name--
A
————▶▶

When two or more items are in a stack below the line, all items are optional. In the following example, you can choose either *A*, *B*, *C*, or nothing.



mount

Use the **mount** command to complete various Data Protection for VMware Recovery Agent tasks.

The Recovery Agent CLI can be used to mount (**mount add**) and unmount (**mount del**) volumes and disks, and to view a list of mounted volumes (**mount view**). To use the **mount** command, Data Protection for VMware Recovery Agent must be running. Use the `set_connection` command to connect a TDPVMwareShell.exe to the mount application.

Snapshots are mounted or unmounted on the system where Data Protection for VMware Recovery Agent is running.

The **mount** command is supported in command mode. The following command types are available. The appropriate tags and parameters are listed alongside each command type.

add Use this command type to mount a disk or volume of a snapshot to the system where Data Protection for VMware Recovery Agent is running. The following list identifies the tags and parameters for the **add** type:

- **-target** - This tag is required.

Use this tag to specify the following targets:

- **Windows** Virtual volume - only for a partition mount
- **Windows** Reparse point - only for a partition mount
- **Windows** **Linux** iSCSI target

The following examples use the **-target** tag:

- **Windows** In the following example *V:* is the virtual volume mount target:

```
-target "V:"
```

- In the following example a reparse point volume mount target is specified:

```
-target "C:\SNOWBIRD@FASTBACK\SnowbirdK\Snowbird\K\\"
```

- **Windows** **Linux** In the following example an iSCSI target is specified:

```
-target "ISCSI: target=<target_name> initiator=<initiator_name>"
```

- **-rep** - This tag is required.

Use it to specify the Tivoli Storage Manager server that is storing the VMware snapshots, and the Tivoli Storage Manager node that has access to the VMware backups. For example:

```
tsm: ip=<ip/host_name> port=<port_number>
node=<node_name> pass=<node_password>
```

- **-type** - This tag is required. Use it to specify that you want to mount a disk or a partition. The options are:
 - type disk
 - type partition
- **VMname** - This tag is required. Use it to specify the VMware machine name that is source of the snapshot. The specified value is case-sensitive.
- **-disk** - This tag is required. Use it to specify the disk number of the source backed up VMware machine to be mounted.
- **-date** - This tag is required. Use to specify the date of the snapshot that you want to mount. For example :
 - date "2011-Jan-12 22:42:52 AM"
- **-PartitionNumber** - This tag is optional. If the -type is partition, enter the partition number to mount.
- **-ro|-fw** - Use this tag to specify whether the mounted volume is read-only (**-ro**) or fake-write (**-fw**).

The following example shows how to specify the **add** type to mount a disk:

```
mount add -rep "tsm: ip=10.10.10.01 port=1500 node=tsm-ba pass=password"
-target "iscsi: target=test1 initiator=initiator_name" -type disk
-vmname VM-03ENT -disk 1 -date "12/9/2010 10:46:57 AM"
```

In this example, a snapshot of VMware named VM-03ent is located on a Tivoli Storage Manager server with IP 10.10.10.01. Disk number 1 of this snapshot is mounted to the system where Data Protection for VMware Recovery Agent is running.

del Use this command type to dismount one or all mounted backups from the system where Data Protection for VMware Recovery Agent is running. The following list identifies the tags and parameters for the **del** type:

- **-target** - This tag is required. Use this tag to specify the target for dismounting. The target for dismounting can be a virtual volume, reparse point, or iSCSI initiator created using the **mount** command. Use *everything* to dismount all mounted backups.
- **-force** - Use this tag to force an unmount. The default option is not to force an unmount if the target is currently in use.

For example, to force an unmount of a snapshot that is currently mounted at the directory, *c:\gever*, use the following command:

```
mount del -target "c:\gever" -force
```

To dismount a snapshot currently mounted as volume *V:*, use the following command:

```
mount del -target V:
```

To dismount a snapshot currently mounted as an iSCSI initiator, use the following command:

```
mount del -target "ISCSI:<target_name>"
```

dump Use this command type to get a list of all the available backups to mount.

- **-rep** - This tag is required. Use this tag to specify the Tivoli Storage Manager server storing the VMware snapshots, and to specify the Tivoli Storage Manager node that has access to the VMware backups. For example:

```

|         tsm: ip=<IP/host name> port=<PortNumber>
|         node=<NodeName> pass=<NodePassword>
|
|     • -file - This tag is optional. Use this tag to identify a file name to store
|       the dump text. If this tag is not specified, the dump text is printed only
|       to stdout.
|

```

The following examples show how to specify the dump type:

- List all the available backed up virtual machines.


```
mount dump -type TSM -for TSMVE -rep P -request
ListVM [-file <FileNameAndPath>]
```
- List all the available disk snapshots of a VMware.


```
mount dump -type TSM -for TSMVE -rep P -request
ListSnapshots -VMName P [-file <FileNameAndPath>]
```
- List all the available partitions of a disk snapshot.


```
mount dump -type TSM -for TSMVE -rep P -request
ListPartitions -VMName P -disk P -date P [-file <FileNameAndPath>]
```

remove

Use this type to remove the connection to a Tivoli Storage Manager server. There is only one tag for the **remove** type:

-rep - This tag is required. Use this tag to specify the Tivoli Storage Manager server connection to be removed.

In the following example, remove the connection to a Tivoli Storage Manager server (10.10.10.01) using node NodeName:

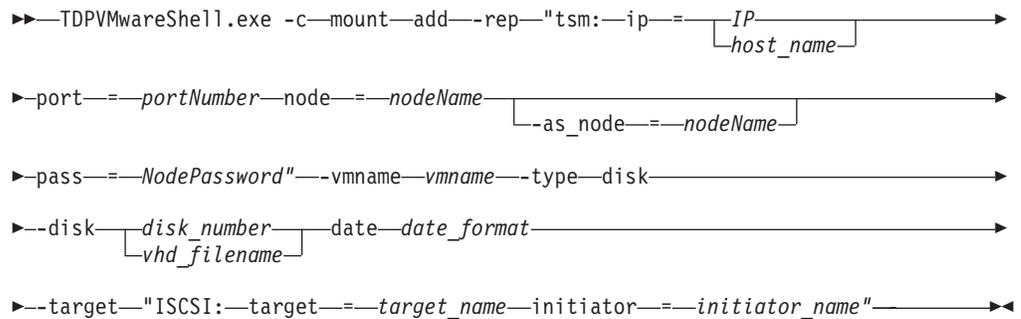
```
mount remove -rep "tsm: NodeName@ip"
```

view Use this type to view a list of all mounted snapshots. This type has no tags. The following example uses the **view** type:

```
mount view
```

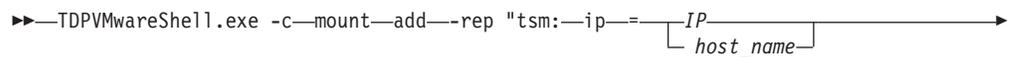
Mounting a disk

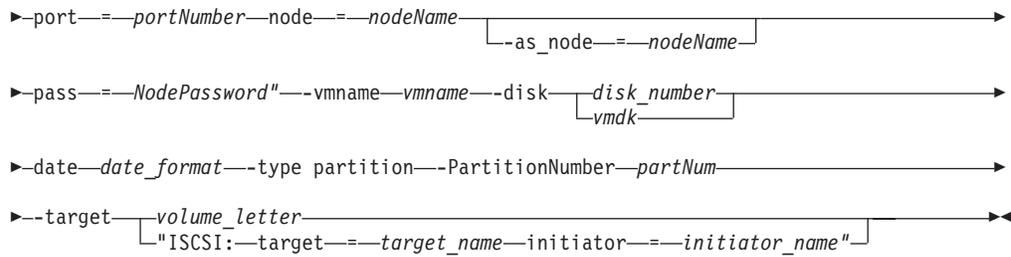
The following syntax diagram is for the command for mounting a disk:



Mounting a partition

The following syntax diagram is for the command for mounting a partition:





set_connection

The **set_connection** command sets the Recovery Agent CLI to work with a specified Data Protection for VMware Recovery Agent.

Use the following format for the **set_connection** command:

```
TDPVMwareShell.exe -c set_connection Command_Tag <hostname or IP address>
```

The following tag can be used with the **set_connection** command:

mount_computer - Use to set the Data Protection for VMware Recovery Agent connection.

In the following example, the Recovery Agent CLI is set to work with Data Protection for VMware Recovery Agent on the *ComputerName* host.

```
set_connection mount_computer ComputerName
```

help

The **help** command displays the help for all of the supported Recovery Agent CLI commands.

Use the following format for the **help** command:

```
TDPVMwareShell.exe -h
```

Recovery Agent command-line interface return codes

Return codes help identify the results of Recovery Agent CLI operations.

Use these return codes to check the status of your Recovery Agent CLI operations.

Table 16. Recovery Agent CLI return codes

Return Code	Value	Description
0	FBC_MSG_MOUNT_SUCCESS	Command submitted successfully to Data Protection for VMware mount.
0	FBC_MSG_DISMOUNT_SUCCESS	Successfully dismantled a snapshot.
0	FBC_MSG_VIEW_SUCCESS	View operation successful.
0	FBC_MSG_DUMP_SUCCESS	Dump operation successful.
0	FBC_MSG_REMOVE_SUCCESS	Remove operation successful.
0	FBC_MSG_IRESTORE_SUBMIT_SUCCESS	Instant restore submitted successfully.
1	FBC_MSG_MOUNT_FAIL	Mount failed (See the mount logs for details).

Table 16. Recovery Agent CLI return codes (continued)

Return Code	Value	Description
2	FBC_MSG_MOUNT_DRIVER_ERROR	Mount driver error.
3	FBC_MSG_VOLUME_LETTER_BUSY	Volume letter or reparse point is in use.
4	FBC_MSG_MOUNT_WRONG_PARAMETERS	Incorrect parameters assigned to the mount command (See the mount logs for details).
5	FBC_MSG_MOUNT_ALREADY_MOUNTED	Job is already mounted on the requested target.
6	FBC_MSG_MOUNT_WRONG_PERMISSIONS	Insufficient permissions.
7	FBC_MSG_MOUNT_NETWORK_DRIVE	Cannot mount on network mapped volume.
8	FBC_MSG_MOUNT_LOCKED_BY_SERVER	Snapshot locked by the server.
9	FBC_MSG_CAN_NOT_CHANGE_REPOSITORY	Cannot change repository.
11	FBC_MSG_DISMOUNT_FAIL	Failed to dismount a mounted snapshot.
13	FBC_MSG_VIEW_FAIL	Retrieving list of virtual volumes failed.
15	FBC_MSG_DUMP_FAIL	Dump command list creation failed.
16	FBC_MSG_CONNECTION_FAILED	Disconnected from Data Protection for VMware mount.
17	FBC_MSG_CONNECTION_TIMEOUT	Operation timed out.
18	FBC_MSG_MOUNT_FAILED_TO_FIND_REPOSITORY	Failed to find a valid repository with snapshots.
19	FBC_MSG_MOUNT_JOB_NOT_FOUND	Failed to find the requested snapshot.
20	FBC_MSG_MOUNT_JOB_FOLDER_NOT_FOUND	Failed to find the requested snapshot data.
22	FBC_MSG_CAN_NOT_REMOVE_REPOSITORY	Cannot remove selected repository.
23	FBC_MSG_REPOSITORY_GOT_MOUNTS	Repository has mounted snapshots.
38	FBC_MSG_MOUNT_NOT_WRITABLE_VOLUME	The mount volume is not writable
39	FBC_MSG_NO_TSM_REPOSITORY	No Tivoli Storage Manager repository was located.
40	FBC_MSG_MOUNT_NOT_ALLOWED_AS_READONLY	Mounting the iSCSI target as read only is not allowed.
41	FBC_MSG_RESOURCE_BUSY_IN_TAPE_MODE	Data Protection for VMware is running in tape mode - media is busy.
42	FBC_MSG_DISK_TYPE_NOT_SUPPORTED	Partition operation not supported for this type of disk.

Table 16. Recovery Agent CLI return codes (continued)

Return Code	Value	Description
43	FBC_MSG_MOUNT_INITIALIZING	The operation failed, Data Protection for VMware mount is currently initializing. Try again later.
44	FBC_MSG_CANNOT_LOCK_SNAPSHOT	The snapshot cannot be protected against expiration during this operation. Refer to documentation for more details.

Appendix. Accessibility features for Tivoli Storage Manager for Virtual Environments

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Manager for Virtual Environments:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices
- User documentation provided in HTML and PDF format. Descriptive text is provided for all documentation images.

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled.

Vendor software

Tivoli Storage Manager for Virtual Environments includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

Related accessibility information

You can view the publications for Tivoli Storage Manager for Virtual Environments in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access these or any of the other documentation PDFs at the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility: <http://www.ibm.com/able>.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.



Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client backup data.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule

A database record that describes the

planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class

See *privilege class*.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

A file or group of files that was archived to server storage.

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B**backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind To associate all versions of a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor

An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX®, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon (CAD)*.

client acceptor daemon (CAD)

See *client acceptor*.

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server

Pertaining to the model of interaction in

distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication

See *data deduplication*.

demand migration

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver

A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-programming interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager, managed server, profile, and subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list

See *include-exclude list*.

execution mode

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSL). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F**file access time**

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G**General Parallel File System**

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state

The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

GPFS™

See *General Parallel File System*.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client

See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup

A backup of a full file system or raw logical volume as a single object.

inactive file system

A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore

See *mailbox restore*.

information lifecycle management (ILM)

GPFS policy-based file management for storage pools and file sets.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local (1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a

storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)

An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A backup of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

managed server

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See *megabyte*.

media server

In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)

(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk

storage capacity and communications volume, 1 000 000 bits.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode

A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N**Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

See *Network Data Management Protocol*.

NetBIOS

See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O**offline volume backup**

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file

A file that contains processing options. On Windows and NetWare systems, the file is called *dsm.opt*. On AIX, UNIX, Linux, and Mac OS X systems, the file is called *dsm.sys*.

originating file system

The file system from which a file was

migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data

mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See *wildcard character*.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and

in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

premigration

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the

amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the `dsmatr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

remote

(1) Pertaining to a system, program, or device that is accessed through a communication line.

(2) For HSM products, pertaining to the origin of migrated files that are being moved.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete

file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script

A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as

they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall

The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization

The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode

A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage

The primary, copy, and active-data storage

pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library

A library device that is used by multiple storage manager servers.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot

An image backup type that consists of a point-in-time view of a volume.

space-managed file

A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management

The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client

A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified

during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool

A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T**tape library**

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can

include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based

systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name

A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V

validate

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or

server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on

Tivoli Storage Manager server storage to their original location.

W**wildcard character**

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

A

accessibility features 107

B

backup
 vmcli command 57
backup vm command 75
books
 See publications

C

commands
 backup vm 75
customer support
 contact xi

D

definitions 113
documentation
 See publications

E

education
 see Tivoli technical training viii
errors 54

F

files
 restore overview 85
 restore task (Linux) 92
 restore task (Windows) 87
fixes, obtaining x
full VM backup 73
 restore using command line 74

G

glossary 113

I

IBM Publications Center v, viii
IBM Support Assistant x
incremental VM backup 74
inquire_config
 vmcli command 61
inquire_detail
 vmcli command 62
installing
 Linux
 language pack 29

instant restore
 overview 85
 task (Linux) 92
 task (Windows) 89
Internet, searching for problem
 resolution ix, x

K

knowledge bases, searching ix

L

LAN environment 77
language pack
 Linux
 installing 29
Linux
 language pack
 installing 29

M

manuals
 See publications
mounting snapshots 77

O

operating systems
 Linux 77
 Windows 77

P

Passport Advantage xi
problem determination 54
 describing problem for IBM Software
 Support xi
 determining business impact for IBM
 Software Support xi
 submitting a problem to IBM
 Software xii
publications
 download v
 order v
 search v
 Tivoli Storage FlashCopy
 Manager viii
 Tivoli Storage Manager v

R

restore
 virtual machine data 77
 vmcli command 59

S

security considerations
 virtual environments 72
set_password
 vmcli command 63
snapshots
 mounting 77
software support
 describing problem for IBM Software
 Support xi
 determining business impact for IBM
 Software Support xi
 submitting a problem xii
Software Support
 contact xi
support contract xi
support information viii
support subscription xi

T

Tivoli technical training viii
training, Tivoli technical viii
troubleshooting 54

V

virtual environments
 security 72
virtual machine data
 restore 77
VM backup 69
vmcli command
 backup 57
 inquire_config 61
 inquire_detail 62
 restore 59
 set_password 63
volumes
 restore overview 85
 restore task (Linux) 92
 restore task (Windows) 87
vStorage backup server
 backing up data 70
 off-host backup 70



Product Number: 5725-A44

Printed in USA

SC27-2898-01

