

IBM Storage Protect for Virtual Environments
8.2.0

Data Protection for VMware User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 247](#).

This edition applies to version 8, release 2 of IBM Storage Protect for Virtual Environments (product number 5725-X00) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2011, 2025.**

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

Contents

About this publication.....	vii
Who should read this publication.....	vii
Publications	vii
What's new.....	ix
Chapter 1. Protection for VMs.....	1
Backup and restore types.....	4
How IBM Storage Protect nodes are used in a virtual environment.....	5
Mount proxy node and data mover node requirements.....	8
How Virtual Volumes are used in a virtual environment.....	10
Tape media guidelines.....	11
Controlling which disks are processed.....	12
VM templates and vApps in a vSphere environment.....	12
Automated client failover.....	13
Out-of-space errors on VMware datastores.....	14
Full VM instant restore environment requirements.....	14
VMware vCenter Server user privilege requirements.....	15
Chapter 2. Managing data with the IBM Storage Protect vSphere Client plug-in.....	19
Getting started.....	19
Available features.....	21
Connecting to the Data Protection for VMware vSphere GUI.....	22
Enabling tagging support.....	22
Setting a data mover node as a tag-based node.....	22
Creating tags in the VMware inventory.....	25
Creating a schedule that is compatible with tagging.....	26
Creating a schedule by using IBM Storage Protect Operations Center.....	26
Creating a schedule by using the DEFINE SCHEDULE command.....	27
Configuring backup policies.....	28
Selecting schedules for backing up virtual machines.....	30
Excluding or including virtual machines from scheduled backup services.....	31
Specifying the retention policy of virtual machine backups.....	32
Selecting a data mover for backing up a virtual machine.....	33
Protecting virtual machine disks by setting the disk protection.....	34
Setting the data consistency of virtual machine backups.....	35
Enabling application protection for a virtual machine.....	36
Managing backup operations for virtual machines.....	37
Managing backup schedules for a vCenter.....	38
Viewing backup schedule history for a vCenter.....	39
Starting an on-demand backup of a virtual machine.....	40
Canceling a backup of a virtual machine.....	42
Viewing the status of backup operations for virtual machines.....	42
Setting the at-risk policy for a virtual machine.....	43
Managing data movers for a vCenter.....	44
Monitoring mount proxies for a vCenter.....	46
Restoring a virtual machine.....	47
Dismounting a virtual machine.....	49
Chapter 3. Getting started with file restore.....	51

Common tasks for restoring files.....	51
File restore prerequisites.....	51
Logging in to restore files.....	54
Restoring files from a virtual machine backup.....	54
Chapter 4. Protection for in-guest applications.....	57
Microsoft Exchange Server data protection in VMware environments	57
Configuring the software for Exchange Server data protection in a VMware environment	57
Managing backups	62
Restoring data.....	66
IBM Storage Protect file space information.....	79
Microsoft SQL Server data protection in VMware environments	79
Configuring the software for SQL Server data protection in a VMware environment.....	80
Managing backups	85
Restoring data.....	88
Sample script for validating full virtual machine backups.....	92
IBM Storage Protect file space information.....	93
Application protection for Active Directory domain controllers.....	93
Chapter 5. Data Protection for VMware commands and options.....	95
dsmc command-line interface.....	95
dsmc commands.....	95
dsmc command options.....	95
vmcli command-line interface.....	97
Backup.....	97
Restore.....	99
Inquire_config.....	104
Inquire_detail.....	106
Set_domain.....	108
Set_option.....	109
Set_password.....	110
Get_password_info.....	112
Start_guest_scan.....	113
Profile parameters.....	115
Recovery Agent command-line interface.....	120
Starting the Recovery Agent command-line interface.....	120
Recovery Agent command-line interface overview.....	120
Chapter 6. Backing up VMware data.....	129
Backing up virtual machine data to IBM Storage Protect.....	129
Setting options for an incremental forever backup schedule.....	131
Backing up virtual machine data to local disks.....	132
Backing up migrated virtual machines.....	133
Backing up organization vDCs to IBM Storage Protect.....	134
Backing up data by disk usage.....	135
Scenario: Including four disks for backup processing.....	136
Scenario: Excluding four disks for backup processing.....	137
Scenario: Separating disks for backup and restore processing.....	138
Backing up virtual machines by domain level.....	139
Scenario: Backing up virtual machines by cluster server.....	140
Scenario: Backing up virtual machines by VMware datastore.....	140
Scenario: Backing up virtual machines by name pattern.....	141
Backing up multiple virtual machines in parallel (optimized backup).....	141
Examples: Backing up multiple virtual machines in parallel.....	143
Backing up virtual machines that host Active Directory controllers.....	146
Specifying a management class to associate objects.....	147
Scenario: Specifying a management class for VMware backups in a vSphere environment.....	147

Scenario: Specifying a management class for VMware control files in a vSphere environment.....	148
Specifying objects to include in backup and restore operations.....	149
Scenario: Specifying objects to include for backup and restore operations in a vSphere environment	149
Chapter 7. Restoring VMware data.....	151
Mounting a virtual machine disk and exporting the volumes.....	151
vSphere environment restore scenario.....	153
Full VM instant restore scenarios.....	154
Full VM instant restore cleanup and repair scenarios.....	155
Full VM instant restore integrity validation scenarios.....	157
Verifying that the Active Directory Domain Controller replicated successfully.....	158
Restoring a virtual disk using multiple sessions.....	160
Restoring virtual disks using a CSV file.....	161
Appendix A. Troubleshooting.....	163
Troubleshooting file restore operations.....	168
Trace options for file restore.....	169
File restore solutions.....	169
File restore operations limitations.....	170
Linux file restore operations limitations.....	172
Windows file restore operations limitations.....	176
VMware attributes.....	176
Accepting self-signed SSL certificates.....	177
Verifying a Data Protection for VMware self-signed web server certificate.....	177
Troubleshooting IBM Storage Protect vSphere Client plug-in problems.....	178
Resolving Platform Services Controller connection problems.....	178
Enabling tracing.....	179
Resolving administrator ID not found messages.....	180
Messages for the IBM Storage Protect vSphere Client plug-in.....	181
Appendix B. IBM Storage Protect recovery agent operations.....	185
Mounting snapshots with the recovery agent.....	185
Restoring files with the recovery agent.....	186
Restoring volumes instantly with the recovery agent.....	187
Restoring volumes instantly from a Windows system with the recovery agent.....	188
Appendix C. Data Protection for VMware vSphere GUI messages.....	191
Appendix D. IBM Storage Protect recovery agent messages.....	213
Appendix E. Accessibility.....	245
Notices.....	247
Glossary.....	251
Index.....	253

About this publication

This publication provides overview, planning, and user instructions for IBM Storage Protect for Virtual Environments: Data Protection for VMware.

Who should read this publication

This publication is intended for administrators and users who are responsible for implementing a backup solution with Data Protection for VMware in one of the supported environments.

In this publication, it is assumed that you have an understanding of the following applications:

- IBM Storage Protect server
- VMware vSphere

Installation, configuration, and upgrade information is documented in the *IBM Storage Protect for Virtual Environments: Data Protection for VMware Installation Guide*.

Publications

The IBM Storage Protect product family includes IBM Storage Protect Plus, IBM Storage Protect for Virtual Environments, IBM Storage Protect for Databases, and several other storage management products from IBM®.

To view IBM product documentation, see [IBM Documentation](#).

What's new in version 8.1.27

IBM Storage Protect for Virtual Environments 8.1.27 introduces updates addressing defects and APARs. For a list of new features and updates in this release and previous version 8 releases, see [Data Protection for VMware updates](#).

If changes were made in the documentation, they are indicated by a vertical bar (|) in the margin.

Chapter 1. IBM Storage Protect for Virtual Environments: Data Protection for VMware overview

IBM Storage Protect for Virtual Environments: Data Protection for VMware provides a comprehensive solution for protecting VMs.

Data Protection for VMware works with the integrated data mover to complete incremental-forever full, and incremental-forever incremental backups of VMs. The data mover node "moves" the data to the IBM Storage Protect server for storage, and for VM image-level restore at a later time. Instant restore is available at the disk volume level and full VM level. In addition, protection of vApps and organization vDCs in a vCloud Director environment is also available.

The data mover is a separately licensed component that contains its own user interfaces and documentation. Familiarity with this product and its documentation is necessary in order to adequately integrate a comprehensive plan for protecting your VMs with Data Protection for VMware. IBM Storage Protect for Virtual Environments for Microsoft Windows includes the data mover data mover features on download package.

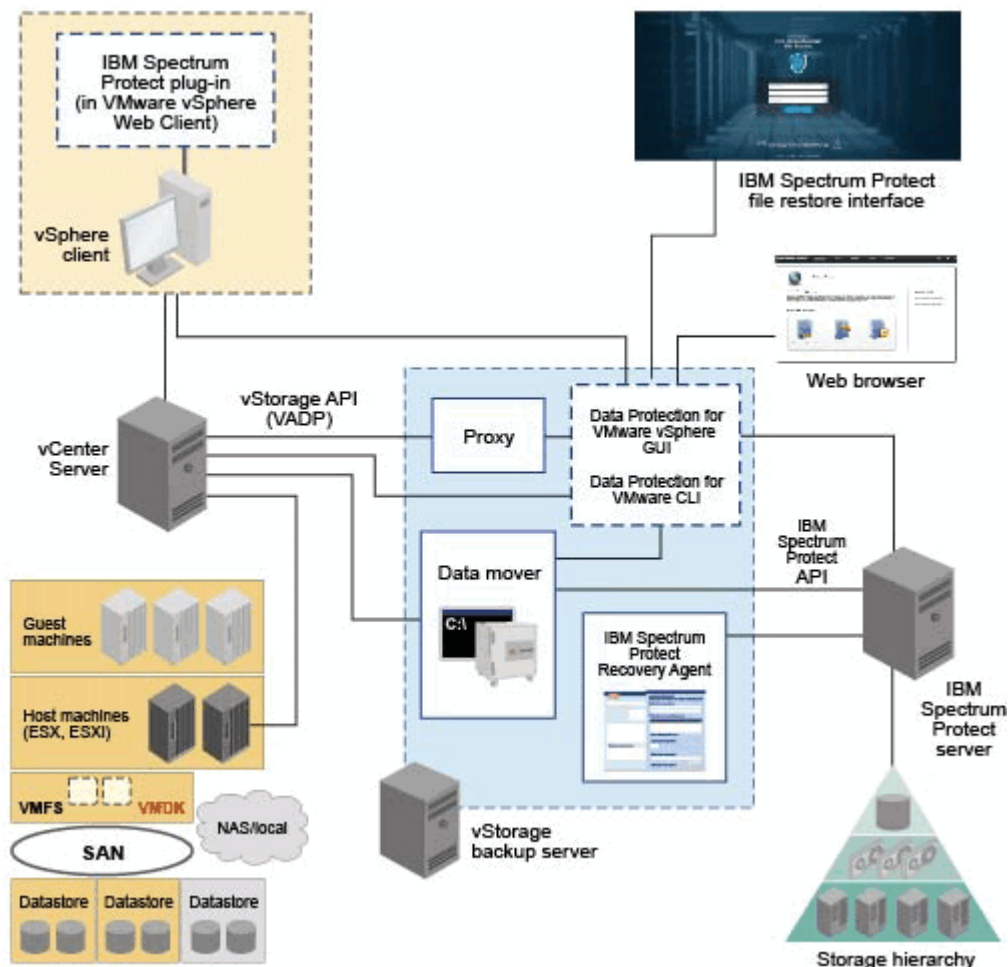


Figure 1. IBM Storage Protect for Virtual Environments system components in a VMware vSphere user environment

Data Protection for VMware provides several components to assist with protecting your VMs.

Data Protection for VMware vSphere GUI

This component is a graphical user interface (GUI) that accesses VM data on the VMware vCenter Server. The content of the GUI is available in two views:

- A web browser view. This view is accessed in a supported web browser by using the URL for the GUI web server host. For example:

```
https://guihost.mycompany.com:9081/TsmVMwareUI/
```

- The IBM Storage Protect vSphere Client plug-in view in the VMware vSphere Web Client. The panels in this view are uniquely designed to integrate within the web client, but data and commands for this view are obtained from the same GUI web server as the other views. The IBM Storage Protect vSphere Client plug-in provides a subset of the functions that are available in the web browser view and some additional functions. Configuration and advanced reporting functions are not offered in this view.

The Data Protection for VMware vSphere GUI can be installed on any system that meets the operating system prerequisites. The Data Protection for VMware vSphere GUI resource requirements are minimal as it does not process I/O data transfers. Installing the Data Protection for VMware vSphere GUI on the vStorage Backup Server is the most common configuration.

For the web-browser view, you can register multiple Data Protection for VMware vSphere GUIs to a single vCenter Server. This scenario reduces the number of datacenters (and their VM guest backups) that are managed by a single VMware Data Protection for VMware vSphere GUI. Each GUI can then manage a subset of the total number of datacenters that are defined on the vCenter Server. For each GUI that is registered to the vCenter Server, one Data Protection for VMware package must be installed on a separate host. To update the managed datacenters, go to **Configuration > Edit IBM Storage Protect Configuration**. In the **GUI Domain** page, reduce the list of datacenters that are managed by the GUI. Managing a subset of all available datacenters reduces the query and processing time that is required by the GUI to complete operations.

When you register multiple Data Protection for VMware vSphere GUIs to a single vCenter Server, the following guidelines apply:

- Each datacenter can be managed by only one installed Data Protection for VMware vSphere GUI.
- A unique VMCLI node name is required for each installed Data Protection for VMware vSphere GUI.
- Using unique data mover node names for each installed Data Protection for VMware vSphere GUI simplifies managing the nodes.

The Data Protection for VMware vSphere GUI must have network connectivity to the following systems:

- vStorage Backup Server
- IBM Storage Protect server
- vCenter Server

In addition, ports for the Derby Database (default 1527) and GUI web server (default 9081) must be available.

Note: The vCenter Server user ID that signs on to the browser view for the Data Protection for VMware vSphere GUI must have sufficient VMware privileges to view content for a datacenter that is managed by the GUI. You must assign this role to a user ID for a specified vCenter Server. The role must be assigned at a vCenter Server level.

IBM Storage Protect file restore GUI

The web-based file restore GUI enables you to restore files from a VMware virtual machine backup without administrator assistance. The GUI is installed automatically when the Data Protection for VMware vSphere GUI is installed. For more information, see [Chapter 3, “Getting started with file restore,” on page 51](#).

IBM Storage Protect recovery agent

This service enables the mounting of any snapshot volume from the IBM Storage Protect server. You can use the iSCSI protocol to access the snapshot from a remote computer. In addition, the recovery agent provides the instant restore function. A volume used in instant restore processing remains available while the restore process proceeds in the background. The recovery agent is accessed with the recovery agent GUI or command-line interface.

If you need to view the snapshot locally with read-only access on the client system, use Data Protection for VMware 8.1.4 or earlier versions.

The recovery agent command-line interface is installed on a Windows system to perform the following tasks from a remote machine:

- Gather information about available restorable data, including lists of:
 - Backed-up VMs
 - Snapshots available for a backed-up machine
 - Partitions available in a specific snapshot

Important: Information about how to complete tasks with the recovery agent GUI is provided in the online help that is installed with the GUI. Click **Help** in any of the GUI windows to open the online help for task assistance.

For detailed information regarding commands, parameters, and return codes, see [“Recovery Agent command-line interface”](#) on page 120.

Data Protection for VMware command-line interface

The Data Protection for VMware CLI is a full-function command-line interface that is installed with the Data Protection for vSphere GUI. You can use it to complete these tasks:

- Initiate a backup of your VMs to the IBM Storage Protect server, or schedule a backup for a later time.
- Initiate a IFFULL recovery of your VMs, VM files, or VM Disks (VMDKs) from the IBM Storage Protect server.
- View configuration information about the backup database and environment.

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

For detailed information regarding available commands, see [“vmcli command-line interface”](#) on page 97.

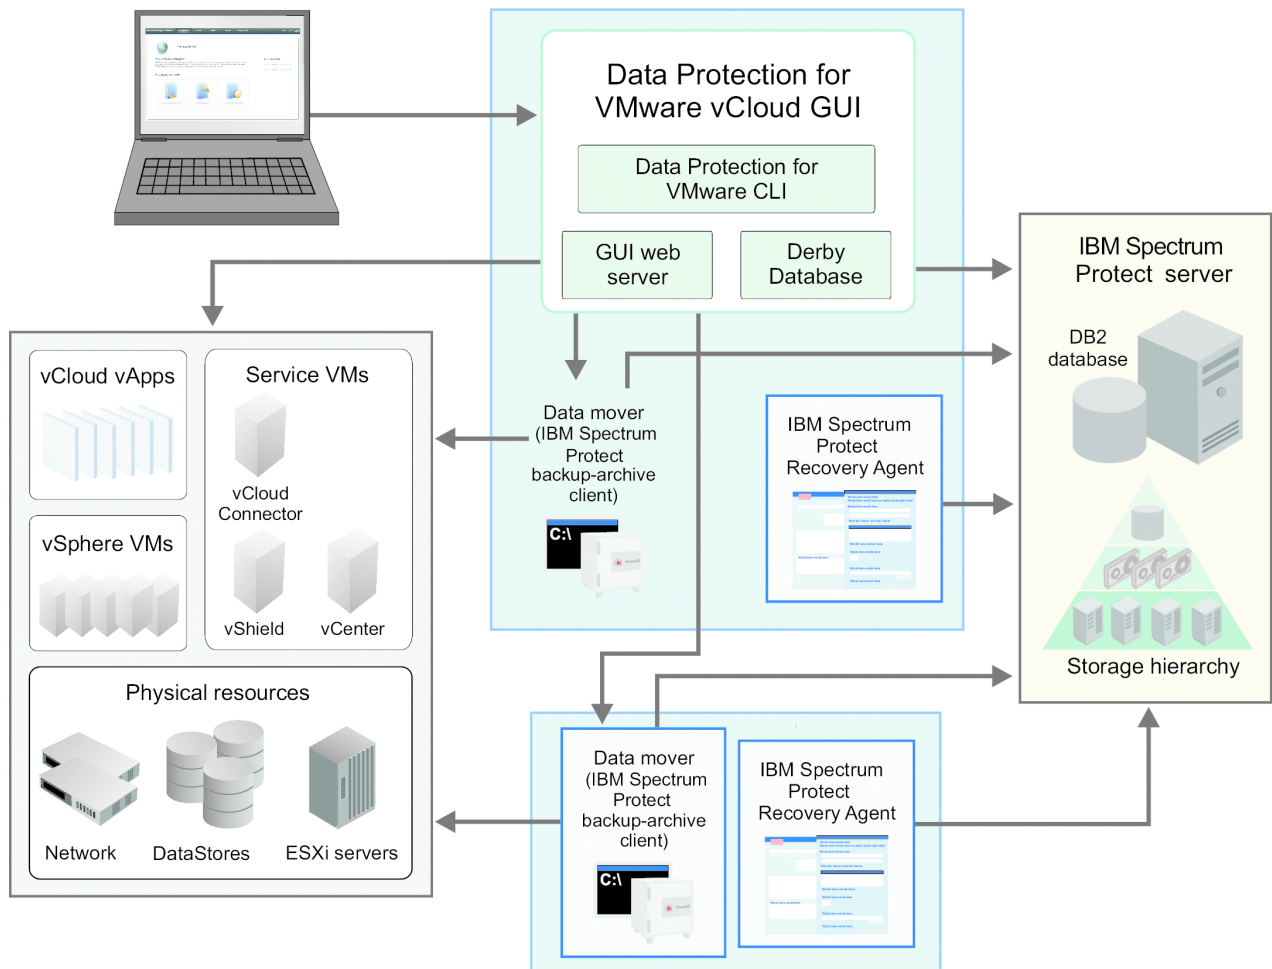


Figure 2. IBM Storage Protect for Virtual Environments system components in a VMware vCloud Director user environment

Backup and restore types

Data Protection for VMware provides multiple backup and restore functions.

Backup types

The following backup types are available:

Incremental-forever incremental backup

Backs up the blocks that changed since the previous backup (full or incremental). The most recent incremental is appended to the previous backup. If a full backup does not exist for this virtual machine (VM), a full backup is automatically performed. As a result, you do not have to verify that a full backup exists.

Incremental-forever full backup

Creates an image of an entire VM. After the full backup is taken, there is no requirement to schedule additional full backups. When full is selected, VM templates that are unchanged since the last backup are also included.

Snapshot differential backup operations are not supported in the VMware environment. You cannot run snapshot differential backup operations of a file system that resides on a NetApp filer on a host where the Data Protection for VMware data mover is also installed.

Restore types

The following restore types are available:

File restore

Use the IBM Storage Protect file restore interface to restore files with a web-based interface. File owners can search, locate, and restore files from a VM backup without administrator assistance.

Fast VM revert

The VM is restored from a persisted snapshot that is on the hardware storage and is available when the restore operation completes. This restore type is available only for VMs that are in a VVOL datastore and can be completed from the IBM Storage Protect vSphere Client plug-in or the command line. For more information, see [How Virtual Volumes are used in a virtual environment](#).

Restore

The VM is restored from a backup that is on the IBM Storage Protect server and is available when the restore operation completes. The entire VM is restored to the state it that it existed in when originally backed up.

Instant restore

The VM is restored from a backup that is on the IBM Storage Protect server and is available when the restore operation completes. The entire VM is restored to the state it that it existed in when originally backed up.

Instant access

A temporary VM is created for verification of the backup data, but the virtual machine is not restored. This restore type requires that you manually dismount the VM when you are finished with it. To dismount the VM, see [Dismounting a virtual machine](#).

How IBM Storage Protect nodes are used in a virtual environment

Data Protection for VMware communicates to VMs during backup, restore, and mount operations through IBM Storage Protect nodes.

A node represents a system on which the data mover, Data Protection for VMware, or other application client is installed. This system is registered to the IBM Storage Protect server. Each node has a unique name (node name) that is used to identify the system to the server. Communication, storage policy, authority, and access to VM data are defined based on a node.

In a Data Protection for VMware vSphere production environment, the most basic node is the data mover node. This node represents a specific data mover that "moves data" from one system to another. In a basic vSphere environment, where VMs are backed up by a single client, the VM data is stored directly under the data mover node.

In some scenarios, several data movers are used to back up a complete virtual environment, such as a VMware datacenter. In this scenario, since the backup work is distributed among multiple data movers, the VM data is stored in a shared node (instead of a specific data mover node). This shared node is called the datacenter node. Thus, in this large system vSphere environment, the data mover nodes store VM data into the datacenter node.

In a large vSphere virtual environment, where multiple data movers and datacenter are operative, a third node is used to communicate among the nodes and the IBM Storage Protect server. This node is the VMCLI node.

A mount proxy node represents the Linux or Windows proxy system that accesses the mounted VM disks through an iSCSI connection. These nodes enable the file systems on the mounted VM disks to be accessible as mount points on the proxy system. You can then retrieve the files by copying them from the mount points to your local disk. Mount proxy nodes are created in pairs and are required by the datacenter node for each Windows or Linux system that serves as a proxy. To increase the number of available mount points, you can configure a datacenter node to have multiple pairs of mount proxy nodes.

Use the Data Protection for VMware vSphere GUI configuration wizard or configuration notebook to set these nodes in a vSphere environment.

Table 1. IBM Storage Protect nodes in a vSphere environment	
Node	Description
vCenter node	The virtual node that represents a vCenter.
datacenter node	The virtual node that maps to a data center. The datacenter nodes hold the data.
VMCLI node	The node that connects the Data Protection for VMware command-line interface to the IBM Storage Protect server and the data mover node.
data mover node	<p>This node performs the data movement.</p> <p>Important: Data Protection for VMware stores sensitive information locally on the data mover, and the data mover might also have direct access to VM storage. Access to the data mover must be protected. Allow only trusted users access to the data mover system.</p>
mount proxy node	This node represents the Linux or Windows proxy system that accesses the mounted VM disks through an iSCSI connection. These nodes enable the file systems on the mounted VM disks to be accessible as mount points.

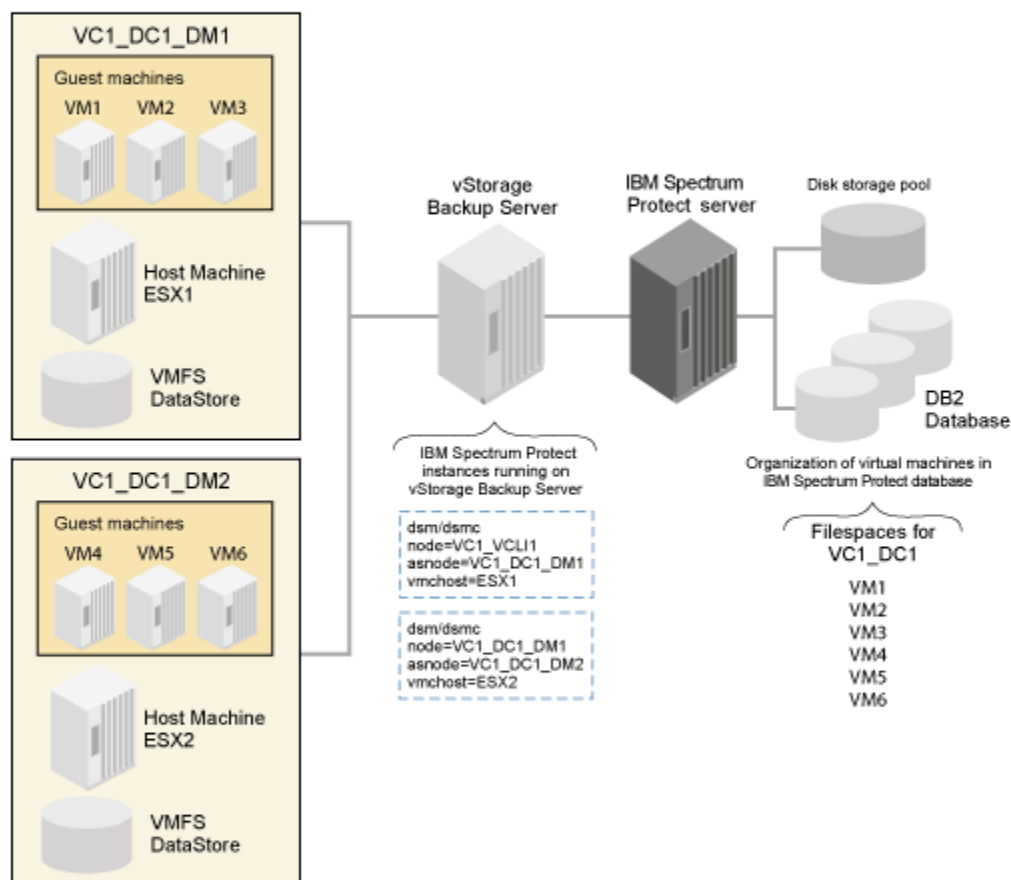


Figure 3. Node relationships and applications in a vSphere production environment that contains one VMware data center and two data move nodes.

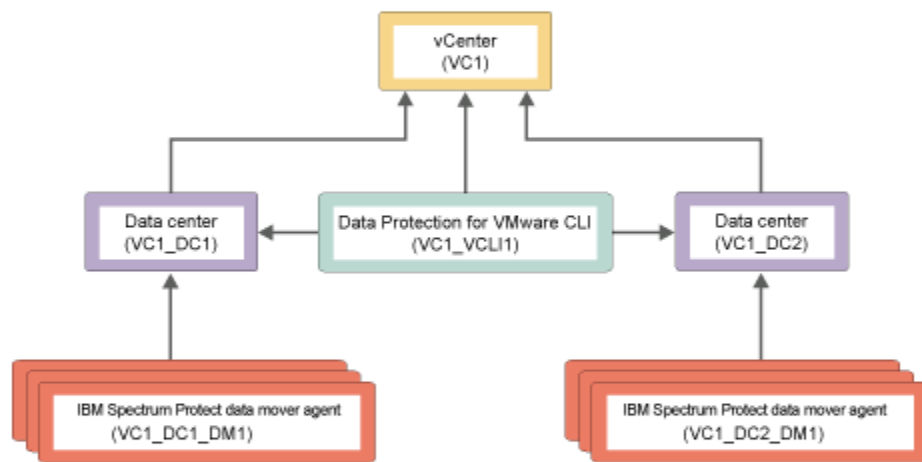


Figure 4. Proxy relationships among the nodes in a vSphere environment that uses two VMware datacenters. The arrows point from the proxy agent node to the proxy target node.

Mount proxy node and data mover node requirements

Operations require specific node types and certain environment settings.

Consider these Data Protection for VMware node requirements before you attempt any tasks:

- Data mover nodes are required for the following operations:
 - IFINCREMENTAL - indicates the incremental-forever incremental backup type.
 - IFFULL - indicates the incremental-forever full backup type.
- Mount proxy nodes are required for the following operations:
 - Full VM instant access
 - Full VM instant restore
 - Mount
- A mount operation accesses a Windows system and a Linux system that function as mount proxy systems. The Windows proxy system also requires the recovery agent to be installed. These two mount proxy nodes function together during a mount operation. Mount proxy nodes are created in pairs and are required by the datacenter node for each Windows or Linux system that serves as a proxy.
- Only one mount proxy node is allowed for each physical or virtual Windows mount proxy system. If you want to use multiple mount proxy node pairs, you must install each Windows mount proxy node on a separate system, along with a recovery agent.
- You cannot mount the backup of a Windows mount proxy node or Linux mount proxy node to itself.
- The following requirements are specific to data movers and mount proxy systems if VMs are in a VVOL datastore:
 - For stability, the data mover and mount proxy should reside on a non-VVOL datastore.
 - For improved performance for file restore for Windows guest VMs, configure the Windows mount proxy system in the same datacenter and the same ESXi host as the guest VMs. This configuration takes advantage of the VMware virtual disk hot add capability.
 - **Linux only:** The credentials for the vCenter must be set on the mount proxy system by using the `dsmc Set Password` command.
 - **Linux only:** The VMCHOST option for the mount proxy must be specified in the `dsm.sys` options file.

The recovery agent is restricted to one node assignment. This node must be a mount proxy node. Although a Windows system might contain multiple data mover nodes, only one proxy mount node is allowed for the recovery agent to use. As a result, operations that use the recovery agent fail when you attempt to connect to a system with a node that is not assigned to the recovery agent.

These examples show types of operations that fail when a node that is not assigned to the recovery agent is used:

Mount operations

When you run a mount operation with the mount proxy node from VMware datacenter DC1, the recovery agent connects to that mount proxy node. Because that connection to the mount proxy node is the only correct connection, the recovery agent does not use another mount operation with any other nodes on that mount proxy system. As a result, the mount operation fails when you use a mount proxy node from VMware datacenter DC2.

Before you attempt a mount operation, you must disable multipathing on the Linux mount proxy system.

Note: Logical Volume Manager (LVM) filtering can block iSCSI connections.

Note: The Linux mount proxy system does not support LVM auto activation.

Instant access or instant restore operations

You attempt to run an instant access or instant restore operation with a mount proxy node from a Windows system that is used as a mount proxy system. A Windows mount proxy system requires the recovery agent to be installed. Because the connection from the recovery agent to the Windows mount proxy node (to run the mount operation) is the only correct connection, an instant access or instant restore operation that attempts to use this mount proxy node (from the same Windows system) fails.

Mount proxy nodes and data mover nodes require proxy authority to the datacenter node. This proxy authority is granted automatically when you set up your nodes with the Data Protection for VMware

vSphere GUI **Configuration Wizard**. However, if you manually set up your mount proxy nodes and data mover nodes, you must grant this proxy authority to the datacenter nodes on the IBM Storage Protect server with the **GRANT PROXYNODE** command. For example:

```
GRANT PROXYNODE TARGET=DC_NODE AGENT=LOCAL_MP_WIN
GRANT PROXYNODE TARGET=DC_NODE AGENT=LOCAL_MP_LNX
```

File sharing security

When you share a mounted virtual machine snapshot, certain security issues can occur that are related to NFS (Linux) and CIFS (Windows) protocols. Review these issues to better understand the security impact when you share a mounted virtual machine snapshot.

When all of the following conditions exist on Linux systems, respective users can access directories on the shared system:

- The mounted volumes that belong to Linux system (B) are shared to a different Linux host (A).
- The Linux host (A) has the same user names as the Linux system (B) that was backed up

For example, *root* user (A) can access all *root* user (B) files, and *tester* (A) can access all of *tester* (B) files. In this situation, the permission group and user are changed to nobody.

This output is an example of access to mounted volumes:

```
esx2vm55:/opt/tivoli/tsm/client/ba/bin # ls -la /CVT/TSM/ESX2VM21/2014-05-22-01_32_53/Volume7
total 19
drwx----- 4 500 500 1024 Apr 28 23:53 .
drwxr-xr-x 8 root root 4096 May 27 22:06 ..
drwxrwxr-x 2 500 500 1024 Apr 28 23:52 RAID_0
drwx----- 2 root root 12288 Apr 28 23:52 lost+found
```

This output is an example of access to shared volumes:

```
[tester1@ESX2VM51 Volume7]$ ls -la
total 19
drwx----- 4 nobody nobody 1024 Apr 28 23:53 .
drwxr-xr-x 8 nobody nobody 4096 May 27 22:06 ..
drwxrwxr-x 2 nobody nobody 1024 Apr 28 23:52 RAID_0
drwx----- 2 nobody nobody 12288 Apr 28 23:52 lost+found
```

Make sure that the correct Linux hostname/IP address or Windows user name is specified. If the correct hostname/IP address or user name is not specified, the share operation fails. This failure is identified by the operating system.

On Windows systems, a user with the same credentials as the backed up Windows virtual machine can access the shared volumes on any Windows system.

How Virtual Volumes are used in a virtual environment

Data Protection for VMware provides data protection for virtual machines (VMs) that are in VMware vSphere Virtual Volumes (VVOL) datastores.

The VVOL infrastructure was introduced in vSphere 6.0 and enables storage management at the VM level rather than the datastore level. For more information about the VVOL model, refer to the VMware product information.

Backing up and restoring VVOL VMs

You can use Data Protection for VMware to complete scheduled backups, on-demand backups, and restore operations for VMs that are in VVOL datastores. This includes the ability complete backup and restore operations for persisted snapshots on the local hardware storage system (local backups) and on the IBM Storage Protect server (server backups).

You can run on-demand backup and restore operations from the following interfaces:

The IBM Storage Protect vSphere Client plug-in

You can use the IBM Storage Protect vSphere Client plug-in to complete on-demand backup and restore operations for local and server backups. The Data Protection for VMware vSphere GUI does not support local backup and restore operations for VMs in VVOL datastores.

The dsmc CLI

Use the **Backup VM** and **Restore VM** commands with the `vmbackuplocation` option

You can run an individual scheduled backup by using the `dsmadmc DEFINE SCHEDULE` command with the `vmbackuplocation` option or run a group of schedules by using the `schedgroup` option. The `schedgroup` option creates a group that contains multiple backup schedules. An example of the use of this option is to group multiple daily local backup schedules with a single IBM Storage Protect server backup schedule.

Maximum retention policy for local backups:

The maximum retention policy is 30 versions for local backups. If there are more than 30 versions, the oldest version is deleted. You can use a management class to specify a retention policy of 1-30 versions.

VMware vCenter Server log in considerations

The vCenter Server ID that the mount proxy uses to log in to the vCenter must have the `VirtualMachine.Provisioning.Clone` privilege.

File restore considerations

The following requirements are specific to Linux® mount proxy systems:

- The credentials for the vCenter must be set on the mount proxy system by using the `dsmc Set Password` command.
- The `VMCHOST` option for the mount proxy must be specified in the `dsm.sys` options file.

Application protection considerations

If you are using Data Protection for VMware with IBM Storage Protect for Mail: Data Protection for Microsoft Exchange Server or IBM Storage Protect for Databases: Data Protection for Microsoft SQL Server for application protection, you can also complete backup and restore operations for VMs that are in VVOL datastores. Database level restore operations for application protection from backups on persisted snapshots (local backups) are not supported. However, database level restores from backups of VMs on VVOL datastores that are stored on the IBM Storage Protect server (server backups) are supported.

For the backup and restore features that are available for each Data Protection for VMware 8.1 release, see [technote 7050037](#).

Tape media guidelines

If your environment contains virtual machine backup data on tape media (such as a tape storage pool or virtual tape library), and the data was either directly stored on tape or migrated to tape over time, consider these guidelines.

Configuration

- Ensure that virtual machine control file data is always on a disk storage pool. You can specify the destination storage pool for virtual machine control file data with the data mover `vmctlmc` option. For more information, see [Vmctlmc](#).
- Use collocation by file space to optimize the tape that contains virtual machine backup data.

When a virtual machine is backed up to the IBM Storage Protect server, each backup is represented as a separate file space on the server. The collocation by file space setting saves data from multiple

IFINCREMENTAL backups of the same virtual machine to the same volume (disk file). When migration to tape occurs, these backups are together on the physical tape.

You can enable collocation at the file space level with the server **COLLOCATE=FILESPEC** parameter. For more information, see [DEFINE STGPPOOL](#).

- Be aware of migration thresholds and how data availability is affected by thresholds. For example, a block in Data Protection for VMware that never changes can be migrated to tape even though the most active backup needs the block.

Recovery

File restore from tape media is not supported. File restore from disk storage is the preferred method.

Consider moving target virtual machine backup data from tape media to disk storage before you attempt a file restore operation. Do not move the virtual machine control data because this data should already be in separate disk storage pool. Also, do not move backup data to the disk storage pool on which the control data resides. If you move backup and control data to the same pool, you will have to complete an IFFULL backup to move the backup and control data to separate pools.

To move backup data, use the server **MOVE NODEDATA** command and ensure that the **FROMstgpool** and **T0stgpool** parameters do not specify pools that contain control data.

Long term retention

Run traditional IFFULL VM backups to tape storage regularly as a solution for long-term storage or tape retention of your data. For example, you can run a IFFULL VM backup to tape monthly as a solution for archive needs.

For additional information related to tape media, see [technote 7021081](#).

Controlling which disks are processed

Set include and exclude statements to control the disks to be processed.

Use Data Protection for VMware in conjunction with the IBM Storage Protect backup-archive client to determine which disks in the VM environment are backed up and restored. VM environments typically contain a combination of system, normal, independent, and raw device mapping (RDM) disks. The ability to extend control granularity to specific disks provides benefits in these situations:

- Recover the disk where the operating system is located in order to replace a corrupted system drive.
- Protect disks in VM environments that use IBM Storage Protect Data Protection applications as guests that contain large database and log files.
- VM configuration information is lost. The VM configuration information is recovered while the disks remain in place.

In previous versions of Data Protection for VMware, a new VM was required whenever a VM restore was performed. If the VM already existed, the restore failed. With this feature, you can restore selected virtual disks but leave the rest of an existing VM intact.

VM templates and vApps in a vSphere environment

Data Protection for VMware supports backing up and restoring VM templates and vApps.

A VM template is a master image of a VM. The template can include an installed guest operating system and a set of applications.

VM templates can be restored to the original VM template, or to an alternative VM template and data store location. Because Data Protection for VMware marks the VM template as one unit, a single file restore of a VM template is not feasible. A single virtual disk cannot be restored, nor can a single virtual disk backup be attached on an ESX host to a target VM.

VMs that are contained in a vApp can be backed up and restored. A vApp is a logical entity that consists of one or more VMs. By using a vApp, you can specify and include all components of a multitier application. A vApp also includes the operational policies and associated service levels of the application contained in the vApp.

The VMs in the vApp are identified in the Data Protection for VMware vSphere GUI as VMs. While you select the VM to back up, you cannot select a vApp. When you restore the VM, the VM is added to one of the following locations:

- If the vApp is present with the original full inventory path, the VM is restored to that location.
- When the original full inventory path is not present or was changed, the VM is restored to the top-level default location on the target ESX host. No containers are created during the restore operation.

When backing up a VM template, and a full backup does not exist for this VM template, the following occurs:

- If the selected backup type is incremental-forever-incremental, and the VM template contains changes, the backup type changes to incremental-forever-full.
- If the selected backup type is incremental-forever-full, this type ensures that the VM template is backed up regardless of whether it contains changes.

Automated client failover

If you backed up data to the IBM Storage Protect server, Data Protection for VMware can automatically fail over to the specified failover server for data recovery when there is an outage on the IBM Storage Protect server.

The server that the IBM Storage Protect data mover node connects to during normal production processes is called the *primary server*. When the primary server and data mover node are set up for node replication, the client data on the primary server can be replicated to another IBM Storage Protect server, which is the *target replication server*. This server also functions as the *failover server*.

When the data mover node logs on during normal operations, connection information for the failover server is automatically sent to the data mover node from the primary server. The failover server information is automatically saved to the client options file on the data mover node. No manual intervention is required by you to add the information for the failover server.

Each time the data mover node logs on to the server, the data mover node attempts to contact the primary server. If the primary server is unavailable, the data mover node automatically fails over to the failover server, according to the failover server information in the client options file. In failover mode, you can restore any replicated client data. When the primary server is online again, the data mover node automatically fails back to the primary server the next time the data mover node connects to a server.

There can be multiple target replication servers. If the client is not able to connect to the source replication server, the client will try to connect to all the target replication servers that have been defined, depending on availability.

Requirements: Before the connection information for the failover server is sent to the client options file, the following requirements must be met:

- The primary server, failover server, and data mover node must be at version 7 (any level) or version 8 (any level).
- The primary and failover servers must be set up for node replication, and all the client nodes must be configured for node replication on the server.
- You must back up VMware data at least once to the primary server.
- Client data on the primary server must be replicated to the failover server at least once.

Restriction: The following restrictions apply to Data Protection for VMware during failover operation:

- Any operations that require data to be stored on the server, such as backup operations, are not available.

- Schedules are not replicated to the failover server. Therefore, schedules are not run while the primary server is unavailable.
- Validation of virtual machine backups is not available.
- All client failover scenarios that use the Data Protection for VMware GUI and the IBM Storage Protect file restore interface are not operational. As a workaround, you can use the **dsmc** command line on the data mover node to restore only data from the failover server (except for VM file level restore, which is not available with the **dsmc** command). You can run a full VM restore, instant restore, or instant access operation by using the **restore vm** command with the **vmrestoretype** option. For more information, see [Vmrestoretype](#).

Windows | Linux Out-of-space errors on VMware datastores

To prevent out-of-space errors during virtual machine backups, you can set a data usage threshold for VMware datastores by using the `vmdatastorethreshold` option.

Use the `vmdatastorethreshold` option to set the threshold percentage of space usage for each VMware datastore of a virtual machine. When you initiate a virtual machine backup, the client checks the data usage of the VMware datastores before the virtual machine snapshot is created. If the threshold is exceeded in any of the VMware datastores, the virtual machine is not backed up.

For example, virtual machine `vm1` spans `datastore1` and `datastore2`. You can issue the following command to ensure that the VMware datastores of a virtual machine are at most 90% full before the virtual machine is backed up:

```
dsmc backup vm vm1 -vmdatastorethreshold=90
```

As a result, the client checks the space usage of both `datastore1` and `datastore2` before the snapshot operation begins. If the space usage of either VMware datastore exceeds the 90% threshold, the backup request for `vm1` is not started.

Requirements:

- Ensure that the threshold is low enough so that the snapshot does not use up all the available space in the VMware datastores. Otherwise, you will run out of space on the VMware datastores and the snapshot will not be created.
- If you use multiple clients that act as data mover nodes, you must add the `vmdatastorethreshold` option to the options file for each data mover.

The client checks the data usage of the VMware datastore that contains the virtual machine disk snapshots. By default, the snapshots are created in the same directory as that of the parent virtual disk (`.vmdk`) file. The client checks the data usage only in the default location.

If you use the `EXCLUDE.VMDISK` option to exclude one or more disks from a backup, the threshold check is still run on these disks. Even though these disks are not backed up, VMware still takes a snapshot of these disks.

Independent disks are not checked during space verification processing because a snapshot of these disks does not use any VMware datastore space.

For more information about the `vmdatastorethreshold` option, see [Vmdatastorethreshold](#).

Windows | Full VM instant restore environment requirements

Review the applications, systems, and versions that are required for full VM instant restore operations.

The following environment requirements must exist before attempting a full VM instant restore operation:

- Full VM instant restore is supported only for IBM Storage Protect data mover 7.1 (or later) Windows 64-bit and Windows vStorage Backup servers.
- Instant access and instant restore capability is supported only for VMware VMs that are hosted on VMware ESXi 5.1 servers, or later versions.

- Full VM instant restore is supported only for disks and virtual tape libraries (VTL). Physical tape storage pools are not supported.
- The IBM Storage Protect recovery agent 7.1 (or later) must be installed on the same system as the data mover 7.1 (or later) data mover system.
- A data mover node that was used for version 7.1.0 instant restore and instant access operations cannot be used for version 8.1.11 instant restore and instant access operations. After you upgrade IBM Storage Protect for Virtual Environments to version 8.1.11, you must create a pair of mount proxy nodes to run instant restore and instant access operations. You can create a mount proxy node pair by using either of the following methods:
 - Go to the **Configuration** window in the Data Protection for VMware vSphere GUI and click **Edit Configuration**. Go to the **Mount Proxy Node Pairs** page and follow the instructions on that page.
 - Follow the steps in [Manually configuring the mount proxy nodes on a remote Windows system](#).
- VMs that were backed up with data mover 6.3 (or later) can be restored by using full VM instant restore.
- With IBM Storage Protect for Virtual Environments 8.1.8 and higher, you can use either a Windows mount proxy node or a Windows data mover node to run instant restore and instant access operations.
- The data mover system requires the IBM Storage Protect for Virtual Environments 7.1 (or later) license file.
- iSCSI mount (with the recovery agent) is used to expose the VM disks to the ESX as virtual RDMs. Instant access and instant restore operations require an iSCSI software or hardware adapter that is configured on the ESX host that is used for these operations.
- Storage vMotion must be installed and configured on the ESX servers that host the VMs to be used during instant restore operations. Instant access operations (that validate the VM backup data) do not require Storage vMotion.
- Instant access and instant restore operations require vSphere privileges that power on VMs (**Virtualmachine.Interaction.PowerOn**).
- NFS data stores are not supported as a temporary data store or target data store for instant access and instant restore operations. This restriction is due to a VMware limitation as described in the VMware Knowledge Base Article, [Unable to add a physical mode disk mapping \(RDM\) to a virtual machine stored on an NFS datastore](#)

For detailed configuration instructions, see [Configuring your environment for full virtual machine instant restore operations](#).

VMware vCenter Server user privilege requirements

Certain VMware vCenter Server privileges are required to run Data Protection for VMware operations.

vCenter Server privileges required to protect VMware datacenters with the web-browser view for the Data Protection for VMware vSphere GUI

The vCenter Server user ID that signs on to the browser view for the Data Protection for VMware vSphere GUI must have sufficient VMware privileges to view content for a datacenter that is managed by the GUI.

For example, a VMware vSphere environment contains five datacenters. A user, "jenn", has sufficient privileges for only two of those datacenters. As a result, only those two datacenters where sufficient privileges exist are visible to "jenn" in the views. The other three datacenters (where "jenn" does not have privileges) are not visible to the user "jenn".

The VMware vCenter Server defines a set of privileges collectively as a role. A role is applied to an object for a specified user or group to create a privilege. From the VMware vSphere web client, you must create a role with a set of privileges. To create a vCenter Server role for backup and restore operations, use the VMware vSphere Client **Add a Role** function.

If you want to propagate the privileges to all datacenters within the vCenter, specify the vCenter Server and select the **propagate to children** check box. Otherwise, you can limit the permissions if you

assign the role to the required datacenters only with the `propagate to children` check box selected. Enforcement for the browser GUI is at the datacenter level.

The following example shows how to control access to datacenters for two VMware user groups. First, create a role that contains all of the privileges defined in [technote 7047438](#). The set of privileges in this example are identified by the role named "TDPVMwareManage". Group 1 requires access to manage virtual machines for the Primary1_DC and Primary2_DC datacenters. Group 2 requires access to manage virtual machines for the Secondary1_DC and Secondary2_DC datacenters.

For Group 1, assign the "TDPVMwareManage" role to the Primary1_DC and Primary2_DC datacenters. For Group 2, assign the "TDPVMwareManage" role to the Secondary1_DC and Secondary2_DC datacenters.

The users in each VMware user group can use the Data Protection for VMware GUI to manage virtual machines in their respective datacenters only.

Tip: When you create a role, consider adding extra privileges to the role that you might need later to complete other tasks on objects.

vCenter Server privileges required to use the data mover

For additional instructions that describe how to set up your environment for full virtual machine instant restore and instant access operations, see the *Configuring Tivoli Storage Manager for Virtual Environments 7.1.x for full virtual machine instant restore operations* technote at <http://www.ibm.com/support/docview.wss?uid=swg21683643>.

The IBM Storage Protect data mover that is installed on the vStorage Backup server (the data mover node) requires the VMCUser and VMCPw options. The VMCUser option specifies the user ID of the vCenter or ESX server that you want to back up, restore, or query. The required privileges that are assigned to this user ID (VMCUser) ensure that the client can run operations on the virtual machine and the VMware environment. This user ID must have the VMware privileges that are described in the above technote.

To create a vCenter Server role for backup and restore operations, use the VMware vSphere Client **Add a Role** function. You must select the `propagate to children` option when you add privileges for this user ID (VMCUser). In addition, consider adding other privileges to this role for tasks other than backup and restore. For the VMCUser option, enforcement is at the top-level object.

vCenter Server privileges required to protect VMware datacenters with the IBM Storage Protect vSphere Client plug-in view for the Data Protection for VMware vSphere GUI

The IBM Storage Protect vSphere Client plug-in requires a set of privileges that are separate from the privileges that are required to sign in to the GUI.

During the installation the following custom privileges are created for the IBM Storage Protect vSphere Client plug-in:

- **Datacenter > IBM Data Protection**
- **Global > Configure IBM Data Protection**

Custom privileges that are required for the IBM Storage Protect vSphere Client plug-in are registered as a separate extension. The privileges extension key is `com.ibm.tsm.tdpvmware.IBMDataProtection.privileges`.

These privileges allow the VMware administrator to enable and disable access to IBM Storage Protect vSphere Client plug-in content. Only users with these custom privileges on the required VMware object can access the IBM Storage Protect vSphere Client plug-in content. One IBM Storage Protect vSphere Client plug-in is registered for each vCenter Server and is shared by all GUI hosts that are configured to support the vCenter Server.

From the VMware vSphere web client, you must create a role for users who can complete data protection functions for virtual machines by using the IBM Storage Protect vSphere Client plug-in. For this role, in

addition to the standard virtual machine administrator role privileges required by the web client, you must specify the **Datacenter > IBM Data Protection** privilege. For each datacenter, assign this role for each user or user group where you want to grant permission for the user to manage virtual machines.

The **Global > IBM Data Protection** privilege is required for the user at the vCenter level. This privilege allows the user to manage, edit, or clear the connection between the vCenter Server and the Data Protection for VMware vSphere GUI web server. Assign this privilege to administrators that are familiar with the Data Protection for VMware vSphere GUI that protects their respective vCenter Server. Manage your IBM Storage Protect vSphere Client plug-in connections on the extension **Connections** page.

The following example shows how to control access to datacenters for two user groups. Group 1 requires access to manage virtual machines for the NewYork_DC and Boston_DC datacenters. Group 2 requires access to manage virtual machines for the LosAngeles_DC and SanFrancisco_DC datacenters.

From the VMware vSphere client, create for example the "IBMDDataProtectManage" role, assign the standard virtual machine administrator role privileges and also the **Datacenter > IBM Data Protection** privilege.

For Group 1, assign the "IBMDDataProtectManage" role to the NewYork_DC and Boston_DC datacenters. For Group 2, assign the "IBMDDataProtectManage" role to the LosAngeles_DC and SanFrancisco_DC datacenters.

The users in each group can use the IBM Storage Protect vSphere Client plug-in in the vSphere web client to manage virtual machines in their respective datacenters only.

Issues related to insufficient permissions

When the web browser user does not have sufficient permissions for any datacenter, access to the view is blocked. Instead, the error message GVM2013E is issued to advise that the user is not authorized to access any managed datacenters due to insufficient permissions. Other new messages are also available that inform users of issues that result from insufficient permissions. To resolve any permissions-related issues, make sure that the user role is set up as described in the previous sections. The user role must have all privileges that are identified in the Required privileges vCenter Server user ID and data mover table, and these privileges must be applied at the datacenter level with the propagate to children check box.

When the IBM Storage Protect vSphere Client plug-in user does not have sufficient permissions for a datacenter, the data protection functions for that datacenter and its content are made unavailable in the extension.

When the IBM Storage Protect user ID (specified by the VMCUser option) contains insufficient permissions for a backup and restore operation, the following message is shown:

```
ANS9365E VMware vStorage API error.  
"Permission to perform this operation was denied."
```

When the IBM Storage Protect user ID contains insufficient permissions to view a machine, the following messages are shown:

```
Backup VM command started. Total number of virtual machines to process: 1  
ANS4155E Virtual Machine 'tango' could not be found on VMware server.  
ANS4148E Full VM backup of Virtual Machine 'foxtrot' failed with RC 4390
```

For further information on the use of privileges, see the note on [**vCenter Server privileges required for the Data Protection for VMware vSphere GUI and data mover.**](#)

To retrieve log information through the VMware Virtual Center Server for permission problems, complete these steps:

1. In **vCenter Server Settings**, select **Logging Options** and set **"vCenter Logging to Trivia (Trivia)**.
2. Re-create the permission error.
3. Reset **vCenter Logging** to its previous value prevent recording excessive log information.

4. In **System Logs**, look for the most current vCenter Server log (vpxd-*xyz*.log) and search for the string NoPermission. For example:

```
[2011-04-27 15:15:35.955 03756 verbose 'App'] [VpxVmomi] Invoke error:  
vim.VirtualMachine.createSnapshot session: 92324BE3-CD53-4B5A-B7F5-96C5FAB3F0EE  
Throw: vim.fault.NoPermission
```

This log message indicates that the user ID did not contain sufficient permissions to create a snapshot (createSnapshot).

Chapter 2. Managing data with the IBM Storage Protect vSphere Client plug-in

The IBM Storage Protect vSphere Client plug-in is a VMware vSphere Web Client plug-in that provides a view of the Data Protection for VMware vSphere GUI.

The IBM Storage Protect vSphere Client plug-in is designed to integrate within the VMware vSphere Web Client, but data and commands for this plug-in are obtained from the Data Protection for VMware vSphere GUI web server.

The IBM Storage Protect vSphere Client plug-in provides a subset of the functions that are available in the browser view for the Data Protection for VMware vSphere GUI and some additional functions. Depending on your environment, you can use this plug-in to configure backup policies to fit your backup management needs, such as excluding or including virtual machines (VMs) in scheduled backup services, changing the retention policy of backups, selecting the VM disks you want to protect, setting the data consistency for backups, and providing application protection for VM backups.

You can also use the plug-in to start on-demand backup and restore operations and to view the most recent backup information for all VMs that are in a vSphere object. This information includes identification of VMs that are at risk of being unprotected because the VM has never been backed up or a backup did not occur in the time interval that is set in the at-risk policy.

Getting started

Learn about the tasks for installing, setting up, and using the IBM Storage Protect vSphere Client plug-in to manage data protection for your VMware datacenter.

You can also find helpful information to get started on the **Getting Started** tab for the IBM Storage Protect vSphere Client plug-in. This tab walks you through the initial configuration of the plug-in, including providing guidance for setting up the plug-in for data protection and a short video that shows you how to configure backup policies for virtual machines (VMs).

This tab also provides information on where to go in the interface to complete common tasks such as modifying backup policies, determining if VMs are at risk, running on-demand backups of VMs, and restoring VMs.

To open the **Getting Started** tab, click **IBM Storage Protect** in the vSphere Web Client object navigator.

Table 2. Roadmap of installation, set up, and management tasks for the IBM Storage Protect vSphere Client plug-in		
Task	Description	Learn more
Install the IBM Storage Protect vSphere Client plug-in	To install the IBM Storage Protect vSphere Client plug-in, select Register as a vSphere Web Client plug-in if you are installing by using the installation wizard. If you are installing in silent mode, use the REGISTER_PLUGIN	IBM Storage Protect vSphere Client plug-in Installing the Data Protection for VMware components
Advanced Installation.	Use this scenario when you want to install a data mover (mount proxy), recovery agent, and required support packages on this system.	Installing the Data Protection for VMware components

Table 2. Roadmap of installation, set up, and management tasks for the IBM Storage Protect vSphere Client plug-in (continued)

Task	Description	Learn more
Configure the information that is required for the IBM Storage Protect vSphere Client plug-in	When the installation wizard completes, the configuration wizard opens. Follow the instructions in the wizard to complete the configuration.	Configuring a new installation with the wizard
Assign privileges for the IBM Storage Protect vSphere Client plug-in to roles	During installation, the custom privileges are created for the IBM Storage Protect vSphere Client plug-in. You must assign these privileges to roles for VMware administrators and users.	VMware vCenter Server user privilege requirements
Connect to the Data Protection for VMware vSphere GUI	The IBM Storage Protect vSphere Client plug-in relies on back-end services that are provided by the Data Protection for VMware vSphere GUI that has been preconfigured for a vCenter. To enable the extension for a vCenter, you must first create a connection to the web GUI for that vCenter.	Connecting to the Data Protection for VMware vSphere GUI
Enable tagging support and configure backup policies	<p>You can use the IBM Storage Protect vSphere Client plug-in to change backup policies such as excluding virtual machines (VMs) from scheduled backup services or changing the retention policy of the VM backups.</p> <p>To use this feature, you must enable support for VMware tagging. You can enable support for tagging from the IBM Storage Protect vSphere Client plug-in or from a tool such as vSphere PowerCLI version 5.5 R2 or later.</p>	Enabling tagging support “Configuring backup policies” on page 28
Ensure that backup schedules are compatible with tagging	<p>Schedules are created by the IBM Storage Protect server administrator to automatically back up virtual machines regularly. To enable schedules to be used with the IBM Storage Protect vSphere Client plug-in, the server administrator must create schedules that are compatible with tagging.</p> <p>The server administrator can create tagging-compatible schedules by using the IBM Storage Protect server DEFINE SCHEDULE command or IBM Storage Protect Operations Center 8.1 or later.</p>	“Creating a schedule that is compatible with tagging” on page 26

Table 2. Roadmap of installation, set up, and management tasks for the IBM Storage Protect vSphere Client plug-in (continued)

Task	Description	Learn more
Manage data protection	Use the IBM Storage Protect vSphere Client plug-in to manage data protection tasks for your VMware datacenter.	“Managing backup operations for virtual machines” on page 37 “Restoring a virtual machine” on page 47
Troubleshooting	Learn how to resolve issues such as Platform Services Controller connection problems, enable tracing, and get more details about IBM Storage Protect vSphere Client plug-in extension messages.	“Troubleshooting IBM Storage Protect vSphere Client plug-in problems” on page 178

Available features

The features that are available in the IBM Storage Protect vSphere Client plug-in depend on the version of VMware vSphere that you are using.

If you are using VMware vSphere 6.0 or later, restore, backup, and data protection tagging functions are available.

The following are a few of the features that are available if you are using VMware vSphere 6.0 or later.

Server and local backups

Depending on your vSphere environment, you can back up virtual machines (VMs) to the IBM Storage Protect server (server backup); to a persisted snapshot on the hardware storage (local backup); or to both locations.

Local backups are available only for VMs that are stored in a VMware virtual volume (VVOL) datastore. If any virtual disk of the VM is not in a VVOL datastore, a local backup is not allowed.

For VMs that are in datastore types other than VVOL, only server backups are available.

Local backups are always a full VM image snapshot, even if incremental backups are configured for the server. Because no network data movement is needed for local snapshots, backup operations can be faster than server backup operations.

Tip: To ensure full protection, do not use local backups exclusively. If you use only local backups, you might encounter issues such as snapshots that are inadvertently deleted or the VM is corrupted or deleted and the snapshots are not accessible.

Schedules and schedule groups

Schedules are used to automatically back up VMs regularly. You assign schedules to the inventory objects in the VMware vSphere Web Client to back up the VMs that are in that object.

Schedules contain the following key attributes:

- The scheduled start time
- The frequency that the schedule is run
- Specification of the `-domain.vmfull=Schedule-Tag` option (and no other domain-level options)

Schedules must be compatible with tagging to be used with the IBM Storage Protect vSphere Client plug-in.

Multiple schedules can be placed in a schedule group. You can then assign the schedule group to an object in the vSphere Web client rather than an individual schedule. An example of the use of a schedule group is to group multiple daily local backup schedules with a single server backup.

Schedules and schedule groups are created by the IBM Storage Protect server administrator. For information about how to create schedules that can be used with the IBM Storage Protect vSphere Client plug-in, see [“Creating a schedule that is compatible with tagging”](#) on page 26.

Connecting to the Data Protection for VMware vSphere GUI

The IBM Storage Protect vSphere Client plug-in relies on back-end services that are provided by the Data Protection for VMware vSphere GUI that has been preconfigured for a vCenter. To use the extension, you must create a connection to the host where the Data Protection for VMware vSphere GUI is installed.

Procedure

To create a connection to the Data Protection for VMware vSphere GUI host:

1. In the vSphere Web Client object navigator, click **IBM Storage Protect**.
2. Click the **ManageConfigure** tab.

The vCenters that you can manage by using the IBM Storage Protect vSphere Client plug-in are shown on the **Connections** page.

3. Select a vCenter, and then click the **Edit** icon.
4. Enter the host name or IP address and port for the Data Protection for VMware vSphere GUI server, and then click **Save**.

Results

If the connection is successful, **Verified Connection** is displayed in the **Connection Status** column for the vCenter.

Note: VMware vSphere GUI host configuration of Flash and HTML interfaces is managed separately. If you make configuration changes in either of these user interfaces, you must update the configuration changes in the other GUI.

Enabling tagging support

IBM Storage Protect uses VMware vSphere tags to establish backup policies for managing protection of virtual machines.

These policies are described in [“Configuring backup policies”](#) on page 28. However, before you can configure backup policies, you must enable tagging support.

Setting a data mover node as a tag-based node

When tagging support is enabled on a data mover node, administrators can apply data protection tags to VMware inventory objects such as host clusters, datacenters, hosts, resource pools, virtual machines, and folders (Host and Cluster folders and VM and Template folders).

Before you begin

Ensure that the following requirements are met:

- VMware vCenter Server must be at Version 6.0 Update 1 or later.
- In order for the Data Protection for VMware vSphere GUI to function correctly with tagging support, ensure that the following requirements are met during the installation of the GUI:
 - At least one data mover and the Data Protection for VMware vSphere GUI must be installed on the same server. This data mover node must be configured so that the vCenter server credentials are

saved. You can save the credentials by running the configuration wizard to save the data mover node password, or by using the **dsmc set password** command on the data mover command line.

If you use other data movers, running on virtual machines or physical machines as additional data movers, you can install them on other servers. For tagging support, all these data movers must also be configured with the `VMTAGDATAMOVER YES` option. These additional data movers do not require the Data Protection for VMware vSphere GUI to be installed on the same server in order for them to work correctly as tag-based data movers.

– Linux

For Linux data movers, ensure that you specify the data mover installation directory and the JavaTM shared library `libjvm.so` in the `LD_LIBRARY_PATH` environment variable. The path to `libjvm.so` is used for tagging support when you enable the `vmtagdatamover` option on the data mover. Starting with 8.1.8, a new script (`spve.sh`) has been added to `/etc/profile.d`. This will correctly set the `LD_LIBRARY_PATH` for the following applications: `dsmc`, `dsmcad` and `dsmj`. This should cover the `libjvm.so` as well. If you see errors with `LD_LIBRARY_PATH`, follow the manual instructions:

1. For IBM Java:

```
export LD_LIBRARY_PATH=/opt/tivoli/tsm/client/ba/bin:$JAVA_HOME/jre/bin/classic
```

For Oracle Java:

```
export LD_LIBRARY_PATH=/opt/tivoli/tsm/client/ba/bin:$JAVA_HOME/jre/lib/amd64/server
```

2. To configure the Client Acceptor Service and Data Mover Scheduler Service to act as a vStorage Backup Server, set the following environment variable in the `/etc/init.d/dsmcad` file:

```
export LD_LIBRARY_PATH=/opt/tivoli/tsm/client/ba/bin
```

Note: On Linux operating systems, the Data Protection for VMware vSphere GUI must be installed by using the default user name (`tdpvmmware`).

- On UNIX and Linux clients, the existing passwords in the `TSM.PWD` files are migrated to the new password store in the same location. For root users, the default location for the password store is `/etc/adsm`. For non-root users, the location of the password store is specified by the `passworddir` option.

The `TSM.PWD` file is deleted after the migration.

Note: For further information on the use of privileges required for working with tagging, see [Installing the Data Protection for VMware components](#)

About this task

You can use data protection tags to configure the backup policy of virtual machines in VMware inventory objects. These data protection tags are presented as settings that can be changed in the IBM Storage Protect vSphere Client plug-in.

Procedure

- Use one of the following methods:

Option	Description
To configure a data mover node using the vSphere plug-in GUI	<ol style="list-style-type: none">1. From the vSphere plug-in, select IBM Storage Protect .2. In the Configure tab, select Data Movers.3. In the Add Data Mover panel, select a datacenter from the drop-down menu.4. Accept defaults, or edit settings for Data mover name, Data mover host name, vCenter User and vCenter password.

Option	Description
	<p>5. Click Add when settings are complete.</p> <p>For further details, see the topic, Setting up data mover nodes with the vSphere plug-in GUI in the Data Protection for VMware vSphere GUI Installation Guide.</p>
<p>To configure a new data mover for tagging support on Windows or Linux by using the Data Protection for VMware vSphere GUI</p>	<ol style="list-style-type: none"> 1. On the system where the Data Protection for VMware vSphere GUI is installed, start the GUI by opening a web browser and entering the GUI web server address. For example: <pre>https://<GUI web server address>:9081/TsmVMwareUI/</pre> 2. Log on with the vCenter user ID and password. 3. Go to the Configuration tab, and select the Edit IBM Storage Protect Configuration action. 4. Go to the Data Mover Nodes page of the configuration notebook. 5. Add a data mover node by completing the following steps: <ol style="list-style-type: none"> a. For the data mover node that you want to set up tagging support for, select Create Services. By default, Tag Based Node is selected to enable the data mover node for tagging support b. To designate the tag-based node as a default data mover node, select Default Data Mover. A default data mover node backs up any new VMs that are added to any container in the datacenter, if the container is already in a protection set. The default data mover also backs up any VMs in the protection set that are not assigned the Data Mover tag. <p>Tip: For Linux systems, if you select a new data mover node as the default tagging node, then remove the <code>vmtagdefaultdatamover</code> line from any other data mover options file that is associated with that datacenter.</p> c. Click OK to save your changes. <p>The <code>vmtagdatamover</code> and <code>vmtagdefaultdatamover</code> (if set) options are added to the data mover options file (<code>dsm.opt</code>).</p>
<p>To configure an existing Windows data mover node for tagging support when the node is on a the same server as the Data Protection for VMware vSphere GUI</p>	<ol style="list-style-type: none"> 1. Complete steps 1-3 in the preceding instructions for configuring a new data mover node for tagging support. 2. On the Data Mover Nodes page, select Tag Based Node for the node that you want to enable tagging support for. 3. Optional: To designate the tag-based node as a default data mover node, select Default Data Mover.
<p>To configure an existing Linux data mover node for tagging support or an existing Windows data mover node that is on a different server than the Data Protection for VMware vSphere GUI</p>	<ol style="list-style-type: none"> 1. Add the <code>vmtagdatamover yes</code> option in the data mover options file (<code>dsm.sys</code> for Linux and <code>dsm.opt</code> for Windows). 2. Optional: To designate the tag-based node as a default data mover node, add the <code>vmtagdefaultdatamover yes</code> or <code>vmtagdefaultdatamover dm_name</code> option to the data mover options file. <p>Tip: For Linux systems, if you select a new data mover node as the default tagging node, then remove the <code>vmtagdefaultdatamover</code> line</p>

Option	Description
	from any other data mover options file that is associated with that datacenter.

Results

After the data mover node is enabled for tagging support, the data mover queries the VMware inventory for tagging information when it runs a backup. The data mover then backs up the virtual machines according to the data protection tags that are set. If the data mover node is not configured for tagging support, any data protection tags are ignored during a backup operation.

Related information

[Vmtagdatamover](#)

[Vmtagdefaultdatamover](#)

[Configuring backup policies](#)

Creating tags in the VMware inventory

IBM Storage Protect tags must be created in the VMware inventory before you can use tagging functions. The tags are created when you use the IBM Storage Protect vSphere Client plug-in or when you run a command on the data mover command line.

About this task

After the data protection tags and categories are created in the VMware inventory, you can use tools such as vSphere PowerCLI Version 5.5 R2 or later to apply these tags to the inventory objects to change their backup policy.

Procedure

Use one of the following methods to create data protection tags and categories in the VMware inventory:

- Use the IBM Storage Protect vSphere Client plug-in to configure backup policies for an inventory object. Changing the backup policy of an inventory object automatically applies the appropriate data protection tags to the object.
- Run the **dsmc set vmtags** command on the data mover node.
You need to run this command only one time to create the tags. You do not need to run the command on every data mover node.

If you are upgrading from a previous version of the data mover software, run the **dsmc set vmtags** command again to create any new tags that are available in the new version of the client.
- From the data mover node, back up a virtual machine in an inventory object with the **vmtagdatamover yes** option in the client options file or as part of the **backup vm** command.
For example: `backup vm testvm -vmtagdatamover=yes`

Results

The data protection settings are created in the VMware inventory. For a list of tags that are created, see [Supported data protection tags](#).

Related tasks

[“Setting a data mover node as a tag-based node” on page 22](#)

When tagging support is enabled on a data mover node, administrators can apply data protection tags to VMware inventory objects such as host clusters, datacenters, hosts, resource pools, virtual machines, and folders (Host and Cluster folders and VM and Template folders).

[“Configuring backup policies” on page 28](#)

You can change the way that backups of your VMware assets are managed, such as excluding or including virtual machines (VMs) in scheduled backup services, changing the retention policy of backups, selecting the VM disks you want to protect, setting the data consistency for backups, and providing application protection for VM backups.

Related information

[Vmtagdatamover](#)

[domain.vmfull](#)

[Set Vmtags](#)

Creating a schedule that is compatible with tagging

Schedules are created by the IBM Storage Protect server administrator to automatically back up virtual machines regularly. To enable schedules to be used with the IBM Storage Protect vSphere Client plug-in, the server administrator must create schedules that are compatible with tagging.

About this task

Schedules are assigned to inventory objects in the vSphere Web Client to back up the virtual machines (VMs) that are in that object. Only schedules with the `-domain.vmfull="Schedule-Tag"` option (and no other domain-level parameters) in the schedule definitions are compatible with tagging support and can be assigned to inventory objects.

Every datacenter configured in the configuration wizard now has an automatically created tagged schedule associated with it in the vSphere plug-in. This daily schedule can be modified within the plug-in to be associated with a different tagged data mover. Other schedule attributes can be configured from the command line or Operation Center.

For ease of use, the recommended method for creating tagging-compatible schedules is to use IBM Storage Protect Operations Center 8.1 or later.

Creating a schedule by using IBM Storage Protect Operations Center

To create a schedule by using a GUI, the IBM Storage Protect server administrator can use IBM Storage Protect Operations Center 8.1 or later.

Before you begin

For information about how to start the Operations Center, see [Opening the Operations Center](#).

Procedure

1. On the Operation Center menu bar, click **Clients Schedules**.
2. On the **Schedules** page, click **+ Schedules**.
The **Create Schedule** wizard opens.
3. Follow the information provided by the **Learn more** link to complete the fields in the wizard. The **Type** and **Subtype** fields must have the following values selected:
Type
Select **Virtual**.
Subtype
Select **VMware**.
4. Associate a data mover with the schedule:
 - a) In the vSphere Web Client, click **IBM Storage Protect > Configure > Schedules**.
 - b) Select a vCenter server from the drop-down list.
 - c) Select the schedule and click **Edit**.
 - d) Select one or more data movers for the schedule.

Results

The new schedule is displayed on the IBM Storage Protect **Configure** > **Schedules** tab of the IBM Storage Protect vSphere Client plug-in with **Yes** in the **Compatible** column.

Related tasks

[“Managing backup schedules for a vCenter” on page 38](#)

To help you manage scheduled backups, you can view the list of IBM Storage Protect schedules that are created for a vCenter.

Creating a schedule by using the DEFINE SCHEDULE command

To create a schedule by using the command line, the IBM Storage Protect server administrator can use the server **DEFINE SCHEDULE** command.

About this task

Tip: For ease of use, if you have IBM Storage Protect Operations Center 8.1 or later and are not using schedule groups, use the Operations Center to create schedules.

For information about how to use this command, including the command parameters and options, see [DEFINE SCHEDULE \(Define a client schedule\)](#).

To be compatible with tagging, the following parameters and options must be included in the schedule definition:

- The `-domain.vmfull="Schedule-Tag"` option (and no other domain-level parameters) must be specified in the option string. The option is case insensitive and must contain no spaces. The quotation marks that enclose the `Schedule-Tag` parameter are optional.
- The schedule must contain the `ACTION=BACKUP` and `SUBACTION=VM` parameters.
- The option string must contain the `-asnodename=datacenter_nodename` option, where the value for the `datacenter` parameter must correspond to the datacenter that is being managed by the IBM Storage Protect vSphere Client plug-in.
- If the `-vmbackuptype=backuptype` option is specified in the option string, the value for the `backuptype` parameter must be `FULLVM` (case insensitive).

Procedure

To create a schedule that is compatible with tagging from the command line:

1. Run the **DEFINE SCHEDULE** command as shown in the following example:

```
define schedule <domain_name>
VMWARE_DAILYFULL_10PM_<datacenter_name>
description="VMware Daily FULL backup for <datacenter_name> starting at 10PM"
action=BACKUP subaction=VM starttime=22:00:00
schedstyle=Classic period=1 perunits=days duration=60
durunits=minutes options='-vmbackuptype=fullvm
-asnodename=<datacenter_nodename>
-mode=IFIncremental
-domain.vmfull="Schedule-Tag"'
```

2. Replace the text in the string `<domain_name>` with the IBM Storage Protect domain that the datacenter and data mover nodes are defined in.
3. Replace schedule name in the string `<datacenter_nodename>` with the relevant datacenter name.
4. Associate a data mover with the schedule by using the following server command:

```
define association domain_name schedule_name data_mover_name
```

Results

The new schedule is displayed on the IBM Storage Protect **Configure** > **Schedules** tab of the IBM Storage Protect vSphere Client plug-in with **Yes** in the **Compatible** column.

Related tasks

[“Managing backup schedules for a vCenter” on page 38](#)

To help you manage scheduled backups, you can view the list of IBM Storage Protect schedules that are created for a vCenter.

Related information

[Domain.vmfull](#)

Creating a schedule group

You can use the `schedgroup` option with the server **DEFINE SCHEDULE** command to create a group that contains multiple schedules. If you select schedule group for an object, all of the schedules in that group are applied to the virtual machines (VMs) that are in the object.

Procedure

To create a schedule group, run the **DEFINE SCHEDULE** command with the `schedgroup` option as shown in the following examples. These examples group two schedules for local backups, `SCHED_A_1` and `SCHED_A_2`, and a schedule for a server backup, `SCHED_A_3`, in to schedule group `GROUP_A`.

Examples of DEFINE SCHEDULE command with the <code>schedgroup</code> option
<pre>define schedule standard SCHED_A_1 Type=Client ACTION=Backup SUBACTION=VM OPTIONS='-vmfulltype=vstor -vmbackuptype=fullvm -vmbackuplocation=server -domain.vmfull="SCHEDULE-TAG" -asnodename=DC_SARTRE_WB -SCHEDGROUP=GROUP_A STARTDate=02/06/2017 STARTTime=06:00:00 SCHEDStyle=Enhanced DAYofweek=ANY</pre>
<pre>define schedule standard SCHED_A_2 Type=Client ACTION=Backup SUBACTION=VM OPTIONS='-vmfulltype=vstor -vmbackuptype=fullvm -vmbackuplocation=local -domain.vmfull="SCHEDULE-TAG" -asnodename=DC_SARTRE_WB -SCHEDGROUP=GROUP_A STARTDate=02/06/2017 STARTTime=12:00:00 SCHEDStyle=Enhanced DAYofweek=ANY</pre>
<pre>define schedule standard SCHED_A_3 Type=Client ACTION=Backup SUBACTION=VM OPTIONS='-vmfulltype=vstor -vmbackuptype=fullvm -vmbackuplocation=local -domain.vmfull="SCHEDULE-TAG" -asnodename=DC_SARTRE_WB -SCHEDGROUP=GROUP_A STARTDate=02/06/2017 STARTTime=18:00:00 SCHEDStyle=Enhanced DAYofweek=ANY</pre>

Tip: Ensure that each schedule in the group can complete before the next schedule is set to start.

For more information about the `schedgroup` and `vmbackuplocation` options, see [Schedgroup](#) and [Vmbackuplocation](#).

Results

The new schedule group is shown with the associated schedules on the IBM Storage Protect **Configure > Schedules** tab of the IBM Storage Protect vSphere Client plug-in.

Configuring backup policies

You can change the way that backups of your VMware assets are managed, such as excluding or including virtual machines (VMs) in scheduled backup services, changing the retention policy of backups, selecting the VM disks you want to protect, setting the data consistency for backups, and providing application protection for VM backups.

Before you begin

Review the information in [Tips for configuring backup policies](#).

About this task

The following VMware inventory objects are the containers that you can use to configure backup policies:

- Datacenter
- Folder (Host and Cluster folders and VM and Template folders)
- Host
- Host cluster
- Resource pool
- Virtual machine

Procedure

1. Navigate to the **Configure Backup Policies** wizard by selecting an inventory object in the vSphere Web Client and completing one of the following actions:
 - Click **Actions > IBM Storage Protect > Manage Data Protection**.
 - Click **Actions > IBM Storage Protect > Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage > IBM Storage Protect > Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure > IBM Storage Protect > Edit**.

This wizard contains a **Server** page to configure server backup policies. If local backups are possible, the wizard also contains a **Local** page to configure local backup policies. For more information about server and local backups, see [“Available features” on page 21](#).

Tip: To view the existing backup policy for an inventory object, select an inventory object and click **Manage > IBM Storage Protect** or **Configure > IBM Storage Protect** depending on the version of vSphere that you are using.

2. Update one or more of the following data protection settings on the **Server** and **Local** pages as applicable. Click one of the links in the **Description** column to learn more about the data protection setting.

Option	Description
Schedule Name	“Selecting schedules for backing up virtual machines” on page 30
Exclude from backup	“Excluding or including virtual machines from scheduled backup services” on page 31
Retention policy	“Specifying the retention policy of virtual machine backups” on page 32
Data mover (VM only)	“Selecting a data mover for backing up a virtual machine” on page 33
Disk protection	“Protecting virtual machine disks by setting the disk protection” on page 34
Data consistency	“Setting the data consistency of virtual machine backups” on page 35
Application protection (VM only)	“Enabling application protection for a virtual machine” on page 36

Tip: If an inheritance icon and object name are displayed in a field, the data protection setting that is shown is inherited from that higher-level inventory object. By changing this setting, you are overriding the inherited property for the current object level and any lower-level objects. For more information about inheritance of data protection settings, see [Inheritance of data protection settings](#).

The data protection settings correspond to data protection tags. For detailed information about the tags, see [Supported data protection tags](#).

3. Click **Finish** after you complete the wizard.

If you want to change all data protection settings back to the inherited states (if any), click **Clear Local Settings**.

Results

After you update the backup policy of an inventory object, data protection tags are assigned to the object. The assigned tags and categories are displayed in the **Tags** portlet in the **Summary** tab of the inventory object.

Selecting schedules for backing up virtual machines

Select a schedule from the **Schedule** field to specify how often and when to automatically back up virtual machines (VMs) in a vSphere inventory object.

About this task

If you select an individual schedule, that schedule is applied to the VMs that are in the object. If you select a schedule group, all of the schedules in that group are applied to the VMs that are in the object. Individual schedules that are in a group are not available for selection.

Schedules can be inherited from a parent inventory object. The **Schedule** field shows the schedule that is used for the inventory object and all child objects. If no schedule is inherited or assigned to the inventory object, a warning message is displayed in the field, and the VMs are not included in any scheduled backups. If you selected multiple inventory objects, a schedule is not shown in the field. You must select an available schedule.

You can override a parent schedule by selecting an available schedule in the **Schedule** field.

When you select a schedule, the **Schedule (IBM Storage Protect)** category and tag are assigned to the inventory object. For more information about this category, see [Supported data protection tags](#)

The value of the tag must match the name of the IBM Storage Protect schedule to be used. All VMs in that container object or in child container objects will be backed up by this schedule. If you do not want to back up certain VMs, you can set the **Excluded** tag on those VMs or on a higher-level container object such as a VM folder.

Procedure

1. Select an inventory object in the vSphere Web Client and complete one of the following actions. You can select a datacenter, folder (Host and Cluster folders and VM and Template folders), host, host cluster, or resource pool.
 - Click **Actions > IBM Storage Protect > Manage Data Protection**.
 - Click **Actions > IBM Storage Protect > Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage > IBM Storage Protect > Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure > IBM Storage Protect > Edit**.
2. On the **Server** page of the **Configure Backup Policies** wizard, select a schedule from the **Schedule** field.

If both the **Server** and **Local** page are provided in the wizard to configure server and local backup options, the selected schedule applies to both server and local backups. The **Schedule** field is not provided on the **Local** page.

Only schedules and schedule groups that are compatible with tagging are shown, and schedule groups must also have an associated data mover.

Auto associate data mover feature:

If one or more schedules do not have an assigned data mover and there is a free data mover available, then the schedule or schedules are presented for selection. If you choose one of these schedules and submit the configuration change, the free data mover is associated with the selected schedule.

A free data mover is one that has no associated schedules and passes verification.

If there is a free data mover on a local system, that data mover is selected. If no local data mover is available, a free data mover from a remote system is selected.

You can view the data mover that was assigned to a schedule on the **IBM Storage Protect > Configure > Schedules** tab in the **Schedule details** section for the selected schedule.

If you do not want to use the assigned data mover for the schedule, select the schedule on the **Schedules** tab and click **Edit**.

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and complete the wizard.

Results

All the VMs in the inventory object and any child objects are protected by the selected schedule, except for any objects that are excluded from scheduled backups.

You can also view the list of IBM Storage Protect schedules that are created for the vCenter. If the schedule belongs to a schedule group, the group is also shown. For more information, see [“Managing backup schedules for a vCenter” on page 38](#).

Related tasks

[“Creating a schedule that is compatible with tagging” on page 26](#)

Schedules are created by the IBM Storage Protect server administrator to automatically back up virtual machines regularly. To enable schedules to be used with the IBM Storage Protect vSphere Client plug-in, the server administrator must create schedules that are compatible with tagging.

[“Excluding or including virtual machines from scheduled backup services” on page 31](#)

You can use the IBM Storage Protect vSphere Client plug-in to include or exclude virtual machines (VMs) from scheduled backup services. You can include or exclude all VMs in an inventory object or individual VMs.

Excluding or including virtual machines from scheduled backup services

You can use the IBM Storage Protect vSphere Client plug-in to include or exclude virtual machines (VMs) from scheduled backup services. You can include or exclude all VMs in an inventory object or individual VMs.

About this task

Typically, the VMs in your VMware datacenter are protected by scheduled backup services with IBM Storage Protect for Virtual Environments: Data Protection for VMware. In some scenarios, you might want to exclude a VM from scheduled backups. For example, you might exclude a VM if it is used only for testing or if it is accessed infrequently.

In other scenarios, you might want to back up only VMs in a certain level of vSphere inventory objects.

The virtual machines must be in a protection set that is protected by a schedule or a schedule group. A protection set consists of the virtual machines in a container that is assigned the Schedule (IBM Storage Protect) tag.

You can set the include and exclude option for server backups, local backups, or both.

Procedure

1. Select an inventory object in the vSphere Web Client and complete one of the following actions. You can select a datacenter, folder (Host and Cluster folders and VM and Template folders), host, host cluster, resource pool, or VM.
 - Click **Actions > IBM Storage Protect > Manage Data Protection**.
 - Click **Actions > IBM Storage Protect > Configure Data Protection**.

- **VMware vSphere 6.0 or earlier:** Click **Manage** > **IBM Storage Protect** > **Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure** > **IBM Storage Protect** > **Edit**.
2. On the **Server** page of the **Configure Backup Policies** wizard, select an item from the **Exclude from backup** list:
 - **Yes** - Excludes the VM from scheduled backups.
 - **No** - Includes the VM in scheduled backups. This selection is the default.

Select **No** to ensure that VMs are included in scheduled backups regardless of inherited settings.

If the selected object is a VM, the **Exclude from backup** setting applies only to the selected VMs.

Tip: If an inheritance icon and object name are displayed in a field, the data protection setting that is shown is inherited from that higher-level inventory object. By changing this setting, you are overriding the inherited property for the current object level and any lower-level objects. For more information about inheritance of data protection settings, see [Inheritance of data protection settings](#).

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and set the **Exclude from backup** option for local backups. You can set the same or different value for the **Exclude from backup** option for server and local backups.

Results

VMs that are in the excluded VMware objects will not be backed up in future scheduled backup operations. However, you can still run an on-demand backup of an excluded object.

Related tasks

[“Starting an on-demand backup of a virtual machine” on page 40](#)

When you start an on-demand backup of a virtual machine (VM), the backup operation begins immediately without waiting for a schedule to run.

Specifying the retention policy of virtual machine backups

You can specify how long to keep a virtual machine (VM) backup or how many versions of the backup to keep on the IBM Storage Protect server or the local hardware storage.

About this task

The retention policy can be either the number of days that backup versions can exist on the server or hardware storage before they expire, or the number of backup versions that exist before they expire. When backup versions expire, they are removed from server storage or hardware storage.

If you do not specify the management class, the retention policy is inherited from a parent object. If no inherited setting exists, the management class that is specified in the vmcc option is used. If the vmcc option is not set, the default retention policy for the datacenter node is used.

The available retention policies are associated with the datacenter, and are created by the IBM Storage Protect server administrator. If more retention policies are required, contact the server administrator.

Procedure

1. Select an inventory object in the vSphere Web Client and complete one of the following actions. You can select a datacenter, folder (Host and Cluster folders and VM and Template folders), host, host cluster, resource pool, or VM.
 - Click **Actions** > **IBM Storage Protect** > **Manage Data Protection**.
 - Click **Actions** > **IBM Storage Protect** > **Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage** > **IBM Storage Protect** > **Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure** > **IBM Storage Protect** > **Edit**.

2. On the **Server** page of the IBM Storage Protect wizard, select a policy from the **Retention policy** list.

If the selected object is a VM, the data protection setting applies only to the selected VM.

Tip: If an inheritance icon and object name are displayed in a field, the data protection setting that is shown is inherited from that higher-level inventory object. By changing this setting, you are overriding the inherited property for the current object level and any lower-level objects. For more information about inheritance of data protection settings, see [Inheritance of data protection settings](#).

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and set the **Retention policy** option for local backups. You can set the same or different value for the **Retention policy** option for server and local backups.

Retention policies that have a version limit greater than 30 or no limit are not shown on the **Local** page.

Results

The retention policy that you set for the VMs in the selected inventory objects will be used for all future backup operations. If the retention policy is changed, the existing backups are rebound to the new retention policy during the next backup.

Related information

[domain.vmfull](#)

Selecting a data mover for backing up a virtual machine

Select a data mover to use for backing up virtual machines from the **Data mover** field. This field is available only for a virtual machine object.

About this task

The data mover is part of the IBM Storage Protect for Virtual Environments: Data Protection for VMware program that backs up VMs to the IBM Storage Protect server or the local hardware storage. The data mover resides on the server where Data Protection for VMware is installed.

The **Data mover** field identifies the data mover that is assigned to a VM or inherited from a parent inventory object. The data mover is inherited from a parent object through the schedule that is assigned to the parent object.

For VMs to be backed up by a schedule, they must be in a container object that belongs to the schedule, and at least one data mover must be associated with the schedule.

If a schedule that is assigned to a container object specifies a single data mover, the VMs inherit the data mover assignment from the container object. However, if the schedule has multiple data mover associations, each VM needs an explicit data mover assignment. Otherwise, the VM will be backed up by the default data mover if one of the associated data movers is configured as the default data mover.

Procedure

1. Select a VM in the vSphere Web Client and complete one of the following actions:
 - Click **Actions** > **IBM Storage Protect** > **Manage Data Protection**.
 - Click **Actions** > **IBM Storage Protect** > **Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage** > **IBM Storage Protect** > **Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure** > **IBM Storage Protect** > **Edit**.
2. On the **Server** page of the **Configure Backup Policies** wizard, select a data mover from the **Data mover** list.

If both the **Server** and **Local** page are provided in the wizard to configure server and local backup options, the selected data mover applies to both server and local backups. The **Data mover** list is not provided on the **Local** page.

All the data movers in the list are associated with the schedule or schedule group that is assigned to the VM and displayed in the **Schedule** field. If no schedule or schedule group is associated with a VM, no data movers are shown.

Tips:

- You can set a data mover for multiple VMs only if you navigated to the **Monitor > IBM Storage Protect** tab, selected multiple VMs that are backed up by the same schedule, and clicked **Actions > Manage Data Protection**.
- If you need to add or remove data movers that are associated with a schedule, click **IBM Storage Protect > Configure > SchedulesIBM Storage Protect > Manage > Schedules**, select the schedule, and click **Edit**.

If you do not assign a data mover to a virtual machine, the data mover is inherited from the parent object. If no inherited setting exists, or the Default Data Mover tag is set or inherited, the virtual machines are backed up by the default data mover that is assigned to a schedule, if any. Otherwise, the virtual machines are not backed up and are identified in the IBM Storage Protect vSphere Client plug-in with the **At Risk** status until a data mover is assigned to the virtual machines.

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and complete the wizard.

4. Optional: To set the data mover selection back to its inherited state (if any), select **Clear** from the **Data mover** field.

Related tasks

[“Managing backup schedules for a vCenter” on page 38](#)

To help you manage scheduled backups, you can view the list of IBM Storage Protect schedules that are created for a vCenter.

Protecting virtual machine disks by setting the disk protection

Select the virtual machine (VM) disks to include in virtual machine backups.

About this task

VM disks are identified by the disk number. For example, in most cases, disk 1 is the system disk.

Complete this procedure if you want to change the default backup behavior, which includes all VM disks in a backup operation. You can also change a non-default behavior that is inherited from a parent object.

Procedure

1. Select an inventory object in the vSphere Web Client and complete one of the following actions. You can select a datacenter, folder (Host and Cluster folders and VM and Template folders), host, host cluster, resource pool, or VM.
 - Click **Actions > IBM Storage Protect > Manage Data Protection**.
 - Click **Actions > IBM Storage Protect > Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage > IBM Storage Protect > Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure > IBM Storage Protect > Edit**.
2. On the **Server** page of the **Configure Backup Policies** wizard, select one of the following settings from the **Disk protection** list:

All disks

Includes all disks in a VM backup.

All disks except disk 1

Includes all disks except disk 1 in a VM backup.

Only disk 1

Includes only disk 1 in a VM backup.

Only disks *n,n,n,...*

Includes a custom set of disks in a VM backup. For example, **Only disks 1,3,5** backs up only disks 1, 3, and 5.

This item is available only if the Disk Backup List category and the Include:*disk number,disk number,...* tag value are set outside of the **IBM Storage Protect** window. For example, the tag Include:1,3,5 includes only disks 1, 3, and 5 in a VM backup.

All disks except disks *n,n,n,...*

Includes all disks except a custom set of disks in a VM backup. For example, **All disks except disks 2,3,4** backs up all disks except for disks 2, 3, and 4.

This item is available only if the Disk Backup List category and the Exclude:*disk number,disk number,...* tag value are set outside of the **Configure Backup Policies** wizard. For example, with the tag, Exclude:2,3,4, all disks are backed up except for disks 2, 3, and 4.

If both the **Server** and **Local** page are provided in the wizard to configure server and local backup options, the selected disk protection setting applies to both server and local backups. The **Disk protection** list is not provided on the **Local** page.

If you do not specify the disks to include or exclude and no inherited setting exists, all virtual machine disks are backed up.

Tip: If an inheritance icon and object name are displayed in a field, the data protection setting that is shown is inherited from that higher-level inventory object. By changing this setting, you are overriding the inherited property for the current object level and any lower-level objects. For more information about inheritance of data protection settings, see [Inheritance of data protection settings](#).

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and complete the wizard.

Setting the data consistency of virtual machine backups

Select the data consistency to achieve for a virtual machine backup operation that fails due to snapshot failure.

About this task

You can set the level of data consistency by specifying the number of snapshot attempts to make, and whether to quiesce the virtual machine file system, including any applications, before attempting the snapshot.

If you do not specify the snapshot attempts and no inherited setting exists, the snapshot attempts that are specified in the `include.vmsnapshotattempts` option are used.

Procedure

1. Select an inventory object in the vSphere Web Client and complete one of the following actions. You can select a datacenter, folder (Host and Cluster folders and VM and Template folders), host, host cluster, resource pool, or VM.
 - Click **Actions > IBM Storage Protect > Manage Data Protection**.
 - Click **Actions > IBM Storage Protect > Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage > IBM Storage Protect > Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure > IBM Storage Protect > Edit**.

2. On the **Server** page of the **Configure Backup Policies** wizard, select one of the following settings in the **Data consistency** list:

Always application consistent

Attempts two file system and Microsoft Windows VSS quiesced snapshots before failing the backup. This selection is the default.

Attempt application consistency

Attempts two quiesced snapshots and, as a final attempt, a nonquiesced, crash-consistent snapshot.

Machine consistent only

Attempts only a nonquiesced snapshot, for VMs that can never complete a quiesced snapshot.

Custom *quiesce,nonquiesce*

Specifies the number of attempts to take a snapshot with quiescing, followed with the number of attempts to take a snapshot without quiescing.

This choice is available only if the Snapshot Attempts category and tag value are set outside of the **IBM Storage Protect** window. In this field, *quiesce* is the number of times to take a snapshot with quiescing, and *nonquiesce* is the number of times to take a snapshot without first quiescing the file system.

For example, with the 2, 2 setting, IBM Storage Protect attempts two quiesced snapshots, and if they fail, attempts two nonquiesced snapshots.

If both the **Server** and **Local** page are provided in the wizard to configure server and local backup options, the selected disk protection setting applies to both server and local backups. The **Data consistency** list is not provided on the **Local** page.

Tip: If an inheritance icon and object name are displayed in a field, the data protection setting that is shown is inherited from that higher-level inventory object. By changing this setting, you are overriding the inherited property for the current object level and any lower-level objects. For more information about inheritance of data protection settings, see [Inheritance of data protection settings](#).

3. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and complete the wizard.

Enabling application protection for a virtual machine

You can use application protection that is provided by IBM Storage Protect.

About this task

When application protection is enabled, IBM Storage Protect notifies virtual machine (VM) applications that a backup is about to occur. This action allows an application to truncate logs and commit transactions before the backup operation begins, so that the application can resume from a consistent state when the backup is completed.

You can enable application protection only on VMs. Ensure that you do not exclude a VM disk (with the **Disk protection** setting) if the disk contains application data that you want to protect.

If you do not enable application protection, the setting in the `include.vmtsmvss` option is used. This setting cannot be inherited.

Procedure

1. Select a VM in the vSphere Web Client and complete one of the following actions:
 - Click **Actions** > **IBM Storage Protect** > **Manage Data Protection**.
 - Click **Actions** > **IBM Storage Protect** > **Configure Data Protection**.
 - **VMware vSphere 6.0 or earlier:** Click **Manage** > **IBM Storage Protect** > **Edit**.
 - **VMware vSphere 6.5 or later:** Click **Configure** > **IBM Storage Protect** > **Edit**.

2. On the **Server** page of the **Configure Backup Policies** wizard, select **Enabled** in the **Application protection** list.

If both the **Server** and **Local** page are provided in the wizard to configure server and local backup options, the selected disk protection setting applies to both server and local backups. The **Application protection** list is not provided on the **Local** page.

3. Optional: If you are backing up a VM that is running a Microsoft SQL Server and want to prevent Microsoft SQL Server logs from being truncated, select **Keep Microsoft SQL Server logs if applicable**.

This option enables the Microsoft SQL Server administrator to manually manage the SQL server logs, so that they can be preserved and be used to restore SQL transactions to a specific checkpoint after the virtual machine is restored. The SQL server administrator must manually back up, and possibly truncate the SQL server logs on the guest virtual machine.

4. If the wizard contains only the **Server** page, click **Finish**.

If the wizard contains the **Server** and **Local** pages, click **Next** and complete the wizard.

5. Ensure that you complete the following configuration steps on each data mover that you are using to back up VMs:

- a) Store the guest VM credentials to Data Protection for VMware by running the following command from the data mover command line:

```
dsmc set password -type=vmguest vm_guest_display_name guest_admin_ID  
guest_admin_pw
```

where *vm_guest_display_name* specifies the name of the guest VM as shown in the VMware vSphere Web Client.

If you use the same credentials to log on to multiple VMs that are enabled for application protection, set the password for all of the VMs by specifying the **allvm** parameter on the following command:

```
dsmc set password -type=vmguest allvm guest_admin_ID guest_admin_pw
```

- b) The command in step “5.a” on page 37 stores the guest virtual machine credentials, which are encrypted on the system that hosts the data mover. Ensure that the following minimum permissions are required for *guest_admin_ID* *guest_admin_pw*:

Backup rights: Microsoft Exchange Server 2013 and 2016: Organization Management permissions (membership in the management role group, Organization Management)

Backup rights: Microsoft SQL Server 2014 and 2016: Organization Management permissions (membership in the management role group, Organization Management)

Managing backup operations for virtual machines

You can use the IBM Storage Protect vSphere Client plug-in to back up your VMware virtual machines (VMs) on the IBM Storage Protect server and to manage and monitor your backups.

About this task

Typically, the VMs in your VMware datacenter are backed up when a schedule is run. Schedules are set up by the IBM Storage Protect server administrator or the VMware administrator to automatically back up virtual machines regularly. You can select a schedule to specify how often and when to automatically back up virtual machines in a vSphere inventory object.

You can also start an on-demand backup of a VM. For example, if you notice that a VM was not backed up recently or if a backup completed with errors, you can start the backup operation again without waiting for the backup to run as scheduled.

You can view the most recent backup information for all VMs that are in a vSphere object. This information includes the backup completion date, duration, and size. This information also includes identification of

VMs that are at risk of being unprotected because the VM has never been backed up or a backup did not occur in the time interval that is set in the at-risk policy.

Related tasks

[“Selecting schedules for backing up virtual machines” on page 30](#)

Select a schedule from the **Schedule** field to specify how often and when to automatically back up virtual machines (VMs) in a vSphere inventory object.

Managing backup schedules for a vCenter

To help you manage scheduled backups, you can view the list of IBM Storage Protect schedules that are created for a vCenter.

About this task

Schedules are set up by the IBM Storage Protect server administrator to automatically back up virtual machines regularly.

A summary of the schedules is displayed in a table that is sortable and filterable on the columns to help you identify and compare the properties of the schedules. For example, you can sort on the **Repeats** column to see which schedules are run the most often.

If there are no schedules in the table, contact your IBM Storage Protect administrator to create a schedule that is compatible with tagging. See [“Creating a schedule that is compatible with tagging” on page 26](#).

Procedure

1. In the vSphere Web Client, click **IBM Storage Protect > Manage > Configure > Schedules**.
2. Select a vCenter server from the drop-down list. The schedules that are defined for the vCenter are shown.
3. To sort the entries in the table, click a column heading.

Information such as the name, start time, frequency, and description of each schedule is shown in the table. You can also see the datacenter that is valid for a schedule. To show more columns, use the scroll bar at the bottom of the table.

The **Compatible** column displays whether a schedule is compatible with the Schedule (IBM Storage Protect) category and tag. Only compatible schedules are supported for tagging and can be assigned to inventory objects in the vSphere Web Client. For information about compatible schedules, see the description for the Schedule tag in [Supported data protection tags](#).

The details of each selected schedule are shown in the **Schedule Details** section of the window.

- The **Data movers** field shows the data movers that are associated with the schedule.
- The **Objects** field shows the inventory objects that are assigned to the schedule. Any virtual machines that are contained in these inventory objects are affected by this schedule.
- To help you diagnose problems with the schedule, the **Options** field shows the IBM Storage Protect options that are defined in the schedule. If necessary, you can validate this information with the IBM Storage Protect administrator.

Tip: Compatible schedules do not use the **Options** field to identify the inventory of VMs to back up.

4. Optional: To modify the data movers that are associated with the schedule to use for backup operations, select a schedule and click **Edit** to open the **Edit Data Movers** window. You can edit only a schedule that is compatible with tagging.
 - a) In the **Data Movers** list, select one or more data movers to use to back up the VMs in the inventory objects that are protected by the schedule.

By default, the selected data movers are assigned to those VMs that either do not have data mover assignments, or are assigned to an invalid data mover for this schedule. Existing data mover assignments that are still valid for the schedule are not overwritten.

Requirement: If you want to use the same data mover in multiple schedules, ensure that the run times of the schedules do not overlap. The data mover can perform backup operations for only one schedule at a time.

- b) In the **Default** column, select the data mover that you want to set as the default for the schedule.

You can also use the `vmtagdefaultdatamover` option with the IBM Storage Protect server

UPDATE SCHEDULE command to set a data mover as the default.

- c) View information about the data mover in the **VMs**, **Data Transmitted**, **Duration**, and **Duration Trend** columns.

- d) If you selected more than one data mover and want to redistribute the workload among the data movers, click **Full rebalance of data movers**. This option assigns the selected data movers to all virtual machines in this schedule.

During a data mover rebalance:

- VMs are sorted by size (storage usage).
- Data movers are assigned to the VMs by size, with the largest VM being assigned to the first data mover on the list, the next largest VM assigned to the next data mover, and so on.
- Existing data mover assignments are overwritten.

For example, if there are only three data movers (DM1, DM2, and DM3) and 10 VMs in a datacenter, the following assignments take place:

- DM1 is assigned to the largest VM.
- DM2 is assigned to the second largest VM.
- DM3 is assigned to the third largest VM.
- DM1 is assigned to the fourth largest VM, and so on.

- e) To save your changes and close the **Edit Data Movers** window, click **OK**.

Related information

[`Vmtagdefaultdatamover`](#)

Viewing backup schedule history for a vCenter

You can view the run history for the backup schedules that are associated with a vCenter. This history includes the dates and times that a schedule ran; the status of the schedule run; and the number of virtual machines (VMs) that were backed up successfully or failed to back up.

About this task

The status of the schedule is based on all data movers that are associated with the schedule. The following status values are possible:

Succeeded

The schedule ran to completion for all data movers.

Failed

The schedule did not run to completion on at least one data mover.

In Progress

The schedule started on all data movers and has not completed.

Pending

The schedule has not started on at least one data mover.

Missed

The schedule failed to start on at least one data mover within the startup window for the schedule.

The number of runs that are shown for a schedule depends on the number of days that are set by the IBM Storage Protect server **SET EVENTRETENTION** command.

Procedure

To view the history of the backup schedules that are defined for a vCenter:

1. In the vSphere Web Client, click **IBM Storage Protect > Monitor > Schedules**.
2. Select a vCenter server from the drop-down list.

You can view the run history for all of the backup schedules that are associated with the vCenter. You can also select a schedule to view the backup status for the VMs that are associated with that schedule in a separate table at the bottom of the **Schedules** page.

Important: Only those VMs for which a backup operation was run are shown when you select a schedule. There might be VMs that are associated with the schedule that are not shown because the backup for the virtual machine did not start.

To see the backup status for all VMs that are associated with an inventory object, select the object and click **Monitor > IBM Storage Protect**. If a backup of a virtual machine did not start, a status of **At Risk** is shown for the virtual machine.

The run history for backup schedules is based on data that is collected from the list of data movers that have a current proxy relationship with the datacenters that are in a vCenter server. If a data mover no longer has a proxy relationship with one of the datacenters, the run history produced by that data mover will not be displayed even if that data mover was previously used to back up VMs.

Starting an on-demand backup of a virtual machine

When you start an on-demand backup of a virtual machine (VM), the backup operation begins immediately without waiting for a schedule to run.

About this task

Typically, the VMs in your VMware datacenter are backed up when a schedule is run. However, you might want to start an on-demand backup if you notice that a VM was not backed up recently or if a backup completed with errors. You can also start an on-demand backup of a VM that is excluded from scheduled backup services.

Tip: Any retention policy settings are observed during an on-demand backup. The retention policy for a VM determines how many backup versions of a VM can exist on the IBM Storage Protect server. Depending on how many backup versions of a VM can exist on the server, running an on-demand backup of a VM might cause older backups to expire. When backup versions expire on the server, they are removed from server storage. For example, if a VM was backed up four times, and only four backup versions can exist on the server, running an on-demand backup causes the oldest backup version to expire.

To check the number of backup versions that can exist on the server, select a VM in the vSphere Web Client and complete one of the following actions:

- **VMware vSphere 6.0 or earlier:** Click **Manage > IBM Storage Protect**.
- **VMware vSphere 6.5 or later:** Click **Configure > IBM Storage Protect**.

You can back up an existing VM by selecting the VM in the vSphere Web Client object inventory and using the **Actions** menu. You can also view and back up an existing VM from the object that contains the VM, such as a datacenter, resource pool, host, or host cluster.

Procedure

To start an on-demand backup of a VM, complete the following steps:

1. Select a VM in the vSphere Web Client object navigator and click **Actions > IBM Storage Protect > Backup**.

You can also select an inventory object that contains the VM or VMs that you want to back up and click **Monitor > IBM Storage Protect**. Complete one of the following actions:

- Right-click one or more VMs, and click **Backup**.
- Select one or more VMs, and click the **Backup** icon.
- Select one or more VMs, and click **Backup** in the **Actions** menu.

2. Complete the following fields in the **Back Up Virtual Machine** window:

Option	Description
Server backup type	<p>Select the type of backup to run:</p> <p>Incremental Backs up the blocks that changed since the previous backup (full or incremental). The most recent incremental is appended to the previous backup. If a full backup does not exist for this VM, a full backup is automatically performed. As a result, you do not have to verify that a full backup exists.</p> <p>Full Creates an image of an entire VM. After the full backup is taken, there is no requirement to schedule additional full backups. When full is selected, VM templates that are unchanged since the last backup are also included.</p>
Data mover	<p>Select the data mover that you want to use to move data to the backup location.</p> <p>Typically, you can accept the default data mover. However, to improve workload balancing in some situations, you might select a data mover that is not heavily used.</p> <p>If there is a data mover assigned to the VM, that data mover is the default. If a data mover is not assigned to the VM or the assigned data mover is no longer valid, the first data mover in the list of data movers that are available for the VM is the default. Available data movers are listed in alphabetical order.</p> <p>Do not select a data mover that is currently in use for a backup or restore operation</p>
Backup destination	<p>Select the destination for the virtual machine backup.</p> <p>Depending on the location of the virtual machine, you can back up to the IBM Storage Protect server, to the local hardware storage, or to both locations.</p>
Disk protection	<p>Select the virtual machine disks to include in backups. The disks are identified by the disk numbers.</p> <p>Custom disk selections are available if the Disk Backup List (IBM Storage Protect) category and tag values are set. See Supported data protection tags.</p>
Data consistency	<p>Select the type of snapshot and retry attempts that are used for backup:</p> <p>Always application consistent Attempts two file system and Microsoft Windows VSS quiesced snapshots before failing the backup.</p> <p>Attempt application consistent Attempts two quiesced snapshots and, as a final attempt, a nonquiesced, crash-consistent snapshot.</p> <p>Machine consistent only Attempts only a nonquiesced snapshot for VMs that can never complete a quiesced snapshot.</p>
Application protection	<p>If a single VM is selected, select Enabled to allow applications such as Microsoft SQL Server and Exchange Server to truncate logs and commit transactions before the backup operation begins. This enables the application to resume from a consistent state when the backup completes.</p>

Option	Description
	<p>Select Keep Microsoft SQL Server logs if applicable if you do not want the logs to be truncated.</p> <p>If a multiple VMs are selected, choose Current value to use the value that is currently set for each virtual machine or select Disabled to disable application protection for all selected virtual machines.</p>

3. Click **OK** to begin the backup operation.
4. Click **Tasks** in the vSphere Web Client object navigator to view the progress of the backup operation.
To see backup tasks, you might have to make sure that tasks for all users are shown. For more information about viewing tasks for all users, refer to the documentation for the version of VMware vSphere that you are using.

Results

After the backup operation is completed, select an inventory object that contains the VM and click **Monitor > IBM Storage Protect** to view the backup information for the VM.

Related tasks

[“Excluding or including virtual machines from scheduled backup services” on page 31](#)

You can use the IBM Storage Protect vSphere Client plug-in to include or exclude virtual machines (VMs) from scheduled backup services. You can include or exclude all VMs in an inventory object or individual VMs.

Canceling a backup of a virtual machine

You can cancel an on-demand or scheduled backup operation that is in progress for a virtual machine (VM).

Procedure

To cancel a backup operation:

1. Click **Tasks** in the vSphere Web Client object navigator.
2. Locate the task for the backup operation that you want to cancel and click the **Cancel Task** icon.

To see backup tasks, you might have to make sure that tasks for all users are shown. For more information about viewing tasks for all users, refer to the documentation for the version of VMware vSphere that you are using.

Results

The backup operation is canceled and **The task was canceled by a user** is shown in the **Status** column for the task.

Viewing the status of backup operations for virtual machines

You can view the most recent backup information for all virtual machines (VMs) that are in a vSphere object. This information includes identification of VMs that are at risk of being unprotected because the VM has never been backed up or a backup did not occur in the time interval that is set in the at-risk policy.

About this task

You can view the backups for virtual machines that are in the following objects in the vSphere Web Client:

- Datacenter
- Folder (host, cluster, and VM)
- Host

- Host Cluster
- Resource Pool

Procedure

To view backup information for the virtual machines that are in an object:

1. Select an inventory object in the vSphere Web Client and click **Monitor > IBM Storage Protect**.

For each VM, information about the most recent backup is shown. For datacenter objects, you can use the **View** list to show all VMs that are in the object, only existing VMs, or only deleted VMs. For all other objects, only existing VMs are shown.

2. For each VM, you can view information such as the backup risk status, completion date, duration, and size.

To view a description of a risk status, hover over the status in the **Risk Status** column.

Related tasks

[“Setting the at-risk policy for a virtual machine” on page 43](#)

Virtual machines (VMs) can be at risk of being unprotected because of failed or missed backup operations. You can set a policy for a VM that specifies if or when the VM is shown as at-risk if a backup operation does not occur in a specified time interval.

Viewing the backup history for virtual machines

About this task

For each task, information such as the backup time, the status of the backup, and the amount of data that was transmitted is shown in the **Backup History** table.

The number of backup tasks that are shown in the **Backup History** table depends on the number of days that are set by the IBM Storage Protect server **SET SUMMARYRETENTION** command.

Procedure

Select a VM from the list on the **Monitor > IBM Storage Protect** tab for an inventory object.

The backup history that is specific to that VM is shown in the **Backup History** table.

You can also select a VM in the object navigator and click **Monitor > IBM Storage Protect** to view the **Backup History** table.

What to do next

To view additional information about the backup tasks for a VM, select the VM, click **Summary**, and view the information in the **Notes** portlet. This summary information includes the information that is in the **Backup History** table and additional information such as the snapshot type and the application protection type if applicable.

Related information

[SET SUMMARYRETENTION](#)

[Supported data protection tags](#)

Setting the at-risk policy for a virtual machine

Virtual machines (VMs) can be at risk of being unprotected because of failed or missed backup operations. You can set a policy for a VM that specifies if or when the VM is shown as at-risk if a backup operation does not occur in a specified time interval.

About this task

By default, the default at-risk policy is set for each VM. You can use the default policy, set a custom policy, or choose to ignore the policy.

You can also change the value for the default policy. This value is applied to all VMs that use the default policy.

Procedure

To change the default at-risk policy, select a custom at-risk policy for selected VMs, or set selected VMs to ignore the at-risk policy, complete the following steps:

1. Select an inventory object that contains the VM in the vSphere Web Client object navigator and click **Monitor > IBM Storage Protect**.
2. Complete one of the following actions:

Option	Description
To set the default at-risk policy	<ol style="list-style-type: none">a. From the Actions menu, click Set Default At-Risk Policy.b. Set the time from the last backup in which the backup operation must occur. The default is 24 hours.
To set a custom at-risk policy	<ol style="list-style-type: none">a. Select the VMs for which you want to set a custom policy.b. From the Actions menu, click Set At-Risk Policy. You can also access the Set At-Risk Policy option by selecting the VMs and clicking the policy icon or right-clicking the selected VMs.c. Click Custom and set the time from the last backup in which the backup operation must occur. The default is 6 hours.
To ignore the at-risk policy	<ol style="list-style-type: none">a. Select the VMs for which you want to suppress at-risk warnings.b. From the Actions menu, click Set At-Risk Policy. You can also access the Set At-Risk Policy option by selecting the VMs and clicking the policy icon or right-clicking the selected VMs.c. Click Ignore.

Results

If the at-risk policy is set to default or custom for a VM, **At Risk** is shown for the VM if a backup does not occur within the time interval that is set by the policy. If the VM has never been backed up, the VM is also considered at-risk and **No Backup** is shown.

If the at-risk policy is set to **Ignore** for a VM, the risk status **Ignored** is shown for the VM regardless of the status of the backup.

Managing data movers for a vCenter

You can manage the data movers that are defined for a vCenter. For example, you can add and remove the association between the data movers and the vCenter and monitor the status of the data movers.

Procedure

1. In the vSphere Web Client, click **IBM Storage Protect > Manage > Configure > Data Movers**.
2. To view data movers, select a vCenter server from the drop-down list. The data movers that are defined for the vCenter are shown.

Optionally, you can view data movers for a specific datacenter or all datacenters that are in the vCenter by clicking **View**.

A summary of the data movers is displayed in a table that is sortable and filterable on the columns. For example, you can click the **Schedule** column to sort the data movers by their associated schedule.

3. Optional: You can complete the following tasks for a data mover:

Action	Steps
Remove data movers	<p>a. Select one or more data movers and click the Remove icon, or select Remove Data Mover from the Actions list.</p> <p>b. Review the information in the Remove Data Movers window and follow the instructions that are shown.</p> <p>Restriction: You cannot remove a data mover that is associated with a schedule. Remove the data mover from all schedules before you attempt to remove the data mover.</p>
Verify data mover status	<p>Select one or more data movers and click the Verify icon, or select Verify Data Mover Status from the Actions list.</p> <p>If the status changes, the value in the Status column is updated. If the status is Failed, you can hover over the status to view the failure message.</p>
Add data movers	<p>a. From the vSphere plug-in, select IBM Storage Protect.</p> <p>b. In the Configure tab, select Data Movers.</p> <p>c. Use the + icon or Actions > Add Data Mover to configure a new Data Mover.</p> <p>d. On the Add Data Mover mover panel, select a datacenter, and specify the data mover name and data mover host name.</p> <p>e. Enter vCenter credentials and click Add. You will be prompted to accept a digital certificate the first time you connect to Add Data Mover operations. For further information on working with digital certificates, see Accepting SSL certificates.</p> <p>To configure a new data mover for tagging support on Windows or Linux:</p> <p>a. Go to the Configuration tab, and select the Edit IBM Storage Protect Configuration action.</p> <p>b. 4. Go to the Data Mover Nodes page of the configuration notebook.</p> <p>c. . For the data mover node that you want to set up tagging support for, select Create Services. By default, Tag Based Node is selected to enable the data mover node for tagging support.</p> <p>For more information, see Configuring a data mover node for tagging support.</p>
View data mover error log	<p>Select to view the data mover error log log file. Click the Download Log button and then select the Save As command from the web client if you want to create a local copy.</p>

Action	Steps
	<p>Note: You may need to enable pop-ups on your browser.</p> <p>Search on the term <code>dsmerror.log</code> on the IBM Documentation for further information.</p>
View data mover schedule log	<p>Select to view the data mover schedule log file. Click the Download Log button and then select the Save As command from the web client if you want to create a local copy.</p> <p>Note: You may need to enable pop-ups on your browser.</p> <p>Search on the term <code>schedule log</code> on the IBM Documentation for further information.</p>

Related concepts

[“How IBM Storage Protect nodes are used in a virtual environment” on page 5](#)

Data Protection for VMware communicates to VMs during backup, restore, and mount operations through IBM Storage Protect nodes.

Monitoring mount proxies for a vCenter

You can view a list of mount proxies that are defined for a vCenter, sort and filter the list, and verify the status of the mount proxies.

Procedure

1. In the vSphere Web Client, click **IBM Storage Protect > Manage > Configure > Mount Proxies**.
2. To view mount proxies, take the following actions:
 - a. Select a vCenter server from the list. The mount proxies that are defined for the vCenter are shown.
 - b. Optionally, to sort the mount proxies, click the columns in the table. For example, you can click the **Status** column to sort the mount proxies by their status: **Verified** or **Failed**.
3. Optional: To verify the status of one or more mount proxies, select the mount proxies and click the **Verify** icon, or select **Verify Mount Proxy Status** in the **Actions** list.

If the status changes, the value in the **Status** column is updated. If the status is **Failed**, you can hover over the status to view the failure message.

Tip:

You can add and remove mount proxies by using the configuration features in the Data Protection for VMware vSphere GUI.

Related concepts

[“How IBM Storage Protect nodes are used in a virtual environment” on page 5](#)

Data Protection for VMware communicates to VMs during backup, restore, and mount operations through IBM Storage Protect nodes.

Restoring a virtual machine

You can restore a virtual machine (VM) from a backup that is on the IBM Storage Protect server or from a persisted snapshot that is on the hardware storage.

About this task

You can restore an existing VM by selecting the VM in the vSphere Web Client object inventory and using the **Actions** menu. You can also view and restore an existing VM from the object that contains the VM, such as a datacenter, resource pool, host, or host cluster. However, if you want to view and restore deleted VMs, you must do so from a datacenter object, which allows you to view all VMs that are in the object, only existing VMs, or only deleted VMs.

Procedure

To restore a VM:

1. Select a VM in the vSphere Web Client object navigator and click **Actions > IBM Storage Protect > Restore**.

You can also select an inventory object that contains the VM and click **Monitor > IBM Storage Protect**. Complete one of the following actions:

- Right-click the VM, and click **Restore**.
- Select the VM, and click the **Restore** icon.
- Select the VM, and click **Restore** in the **Actions** menu.

2. Complete the following pages in the **Restore a Virtual Machine** wizard as applicable. The pages that are provided depend on the options that you select in the wizard.

Option	Description
Select restore point	<p>Use this page to select the backup that you want to use for the restore operation. The VM is restored to the state that it existed for the selected backup.</p> <p>This page also shows the restore options that are available for the VM. These options depend on the location of the backup. The backup can be located on the IBM Storage Protect server, the hardware storage as a persisted snapshot, or both locations.</p>
Select options	<p>Use this page to create a new VM or replace the existing VM with the data from the selected restore point. If you create a new VM, the default VM name is the original name appended with a date and time. You can use this name or enter another name that is not already used by another VM in the datacenter.</p> <p>You can select one of the following restore types. The restore types that are available depend on the whether you are creating or replacing the VM and the restore options that are shown on the Select restore point page.</p> <p>Fast VM revert The VM is restored from a persisted snapshot that is on the hardware storage and is available when the restore operation completes.</p> <p>Restore The VM is restored from a backup that is on the IBM Storage Protect server and is available when the restore operation completes.</p> <p>Instant restore The virtual machine is restored from a backup that is on the IBM Storage Protect server and is available during the restore operation.</p>

Option	Description
	<p>Instant access</p> <p>A temporary virtual machine is created for verification of the backup data, but the virtual machine is not restored.</p> <p>This restore type requires that you manually dismount the VM when you are finished with it. To dismount the VM, see Dismounting a virtual machine.</p> <p>The host that was used when the backup operation was completed is selected by default in the Select the host to restore the virtual machine to list. You can accept the default or select another host. If you selected Fast VM revert, you cannot select a different host.</p>
Select resources	<p>Use this page to restore all disks for the VM and the VM configuration or to restore only selected disks. If you restore to selected disks, you can specify whether you want to restore just the disks, just the configuration, or both for the VM.</p> <p>If you selected Create a new virtual machine on the Select options page, the Restore the virtual machine configuration option is not available.</p>
Select storage	<p>Use this page to select the datastore for the VM. For instant restore operations, you must also select a temporary datastore in the Select temporary datastore list or accept the selected default temporary datastore. Files that are created or updated for the restore operation are saved in the temporary datastore and then copied to primary datastore when the restore operation is complete.</p> <p>If you selected the restore type Restore on the Select options page, you can restore the VM using the same disk format as the source disks or you can select the Thick or Thin format option from the Select the virtual disk format list. The default is Save format as source.</p> <p>If you selected the restore type Instant restore, you can select Thick or Thin from the Select the virtual disk format list. The default is Thick.</p> <p>To view datastores that are compatible with VMware storage policies, select the policy from the Filter by virtual machine storage policy list. The datastores are listed as compatible or incompatible with the selected policy.</p>
Select data mover or Select mount proxy	<p>Use this page to select the data mover or mount proxy that you want to use to for the restore operation.</p>
Ready to complete	<p>Use this page to review the options that you selected in the wizard. Click Finish to start the restore operation.</p>

- Click **Tasks** in the vSphere Web Client object navigator to view the progress of the restore operation. To see restore tasks, you might have to make sure that tasks for all users are shown. For more information about viewing tasks for all users, refer to the documentation for the version of VMware vSphere that you are using.

Results

After the restore operation is complete, the VM is restored in the location that you selected.

Dismounting a virtual machine

Operations such as instant access, instant restore, and file restore require that the associated virtual machine (VM) is mounted. When you are finished using a mounted virtual machine, you can dismount it to clean up the resources on the mount proxy machine that was used for the mount.

About this task

Mounts for instant restore and file restore operations will dismount automatically and usually require a manual dismount only if an error occurs or the mount remains for an extended time. Mounts for instant access operations and general purpose mounts must be dismounted manually.

Procedure

To dismount the VM manually, complete the following steps by operation type:

1. In the vSphere Web Client, click **IBM Storage Protect > Monitor > Maintenance** to view mounted VMs.
2. Review the information for each VM mount, including the operation type and the status (Succeeded or Failed).
3. Depending on the operation type and status, complete the following steps:

Option	Description
General purpose, file restore, or instant access cleanup (Succeeded or Failed status) Instant restore (Succeeded status)	<ol style="list-style-type: none">a. On the Maintenance tab, select a vCenter server from the drop-down list.b. Select the entry for the VM that you want to dismount and click Cleanup. <p>To view the progress of the cleanup by selecting the VM in the vSphere Web Client object navigator and viewing the task and event information on the Monitor tab. For additional troubleshooting information, see task and event information on the Monitor tab for the VM host.</p>
Instant restore (Failed status)	<ol style="list-style-type: none">a. In the vSphere Web Client object navigator, click IBM Storage ProtectConfigure Connections.b. Select the vCenter that contains the VM mount.c. Select the entry for the VM and click the Open icon to open the Data Protection for VMware vSphere GUI.d. Click Restore > Instant Access/Restore Status.e. Review the Action Needed column for the action that is required to dismount the VM. The action might require that you repair the VM. Click the Help icon to open online help for task assistance.

Results

The VM is dismounted and deleted from the associated datacenter.

Chapter 3. Getting started with file restore

To restore files from a web-based interface without administrator assistance, you can use file restore. After the configuration is complete, file owners can search, locate, and restore files.

The web-based interface does not require a file manager application to manually copy files. When you restore a file, you specify a restore point, search or browse to locate the file, and start the restore.

When the configuration is complete, no administrator interaction is required to access or restore files. During the configuration process, the administrator gives the file owner access to the virtual machine that contains their data. File owners access the data with local virtual machine credentials so administrators can monitor file restore resources. File owner permissions do not have to be managed.

You can view demonstration videos that help you learn about the IBM Storage Protect file restore interface. The *Find and Restore Files* and *Monitoring Restores* videos display when you first log in to the file restore interface. Videos are available in English only.

Common tasks for restoring files

Different types of users set up and use the file restore feature. Each user is responsible for a set of tasks.

File owner

The file owner maintains business data such as text documents, spreadsheets, and presentation files.

The file owner completes the following tasks:

- [“Logging in to restore files” on page 54.](#)
- [“Restoring files from a virtual machine backup” on page 54](#)

Administrator

The administrator creates initial software deployments, schedules virtual machine backups to the IBM Storage Protect server, and manages user accounts and permissions in the VMware environment.

The administrator completes the following tasks to set up the environment for file restore:

1. [Enabling the environment for file restore operations](#)
2. [“Backing up virtual machine data to IBM Storage Protect” on page 129](#)
3. Optional: [Setting up file restore operations on Linux](#)

When running Data Protection for VMware in a Linux environment, the file restore feature must be installed on a Windows system to enable the file restore feature.

After the environment is ready for file restore operations, the following optional tasks can be done by the administrator:

- [Modifying options for file restore operations](#)
- [Configuring log activity for file restore operations](#)

File restore prerequisites

Before you restore files with the IBM Storage Protect Data Protection for VMware file restore interface, ensure that your environment meets the minimum prerequisites.

To enable the file restore feature, Data Protection for VMware must be installed on a Windows system.

VMware virtual machine prerequisites

The following prerequisites apply to the VMware virtual machine that contains the files to be restored:

- **Windows | Linux** VMware Tools must be installed on the virtual machine.
- **Windows | Linux** The virtual machine must be running during the file restore operation.
- **Windows** The data mover system must either belong to the same windows domain or be in a domain with a trust relationship with the virtual machine that contains the files to be restored.
- **Windows** When a virtual machine is deleted from a Windows domain and then restored later, the virtual machine must rejoin the domain to ensure the domain trust relationship. Do not attempt a file restore from the virtual machine until the domain trust relationship is restored.
- **Windows** If the user does not own the file to be restored, the Microsoft Windows Restore Files and Directories privilege must be assigned to the user for that virtual machine.
- For further information on Microsoft Windows domain account prerequisites required to use the Data Protection for VMware file restore interface, see [technote 1998066](#).
- **Linux** Local user authentication is required for the virtual machine. Authentication is not available through Windows domain, Lightweight Directory Access Protocol (LDAP), Kerberos, or other network authentication methods.
- **Linux** On a Red Hat Enterprise Linux 6 operating system, the ChallengeResponseAuthentication option in the sshd daemon configuration file (/etc/ssh/sshd_config) must specify YES or be commented out. For example, either of the following statements are valid:

```
ChallengeResponseAuthentication yes
```

```
#ChallengeResponseAuthentication no
```

Restart the sshd daemon after you modify this option.

Data mover prerequisites

The data mover system represents a specific data mover that "moves data" from one system to another.

- **Windows** The data mover system must belong to the same Windows domain as the virtual machine that contains the files to be restored.

Mount proxy prerequisites

The mount proxy system represents the Linux or Windows proxy system that accesses the mounted virtual machine disks through an iSCSI connection. This system enables the file systems on the mounted virtual machine disks to be accessible as restore points to the file restore interface.

- **Linux** Linux operating systems provide a daemon that activates Logical Volume Manager (LVM) volume groups as these groups become available to the system. Set this daemon on the Linux mount proxy system so that LVM volume groups are not activated as they become available to the system. For detailed information about how to set this daemon, see the appropriate Linux documentation.

- **Windows | Linux** The Windows mount proxy system and Linux mount proxy system must be on the same subnet.

The Btrfs file system is removed in Red Hat Enterprise Linux 8. To restore objects that were backed up from a Btrfs file system, use a SUSE Linux Enterprise Server (SLES) based mount proxy node.

Microsoft Windows domain account prerequisites

The following prerequisites apply to Windows domain accounts. The first requirement is to establish a Windows domain user account with local administrative authority over all VMs:

- To perform the necessary tasks to enable file recovery to a virtual machine guest, you need a user account that belongs to a Windows domain and is a local administrator on the mount proxy system . An administrator with this account enters the account credentials in the Data Protection for VMware vSphere GUI configuration wizard or notebook to enable the environment for file restore operations.
- To create a user account with sufficient privileges to use the file restore interface, you can use the Windows Group Policy object to centrally manage a single domain user, allow it to access multiple machines with local administrator credentials, and optionally restrict undesirable actions.

The following steps illustrate how this user account can be created. Complete these steps on a domain controller by using the Active Directory Users and Computers MMC snap-in:

1. Select **Action->New->Groups** and create a new security group named **FR Admins**. The group scope should be set to Global.
2. Create a new domain user account with the user name `fradmin1` and add it to the **FR Admins** security group. You can also add other domain user accounts to the group.
3. To provide more control over the set of computers that `fradmin1` can access, create a new organizational unit
4. From the domain object, select **New->Organizational Unit**, name it as `FR Computers`
5. Populate the `FR Computers` organizational unit with a number of machines. .

Complete the following steps on the domain controller from the Group Policy MMC snap-in:

1. Create a new Group Policy object named `FR Admin GPO`, which will add the administrators in the **FR Admins** group to the local administrator group of the computers associated with the organizational unit to which the Group Policy object is applied.
2. In the Group Policy object, add the account to both the local administrator group and optionally to remote desktop users.
3. Select the `FR Computers` organizational unit and add the newly created Group Policy object.
Note: The Group Policy object could have been associated with the domain itself, but then `fradmin1` would be in the local administrator group of all computers in the domain. Using an explicit organization unit provides additional control.
4. Optionally: use Group Policy Management to restrict undesirable actions on the local machine such as `Deny log on locally` and `Deny log on through Terminal Services`.
5. On the File Restore page of the Data Protection for VMware vSphere GUI configuration wizard or notebook, update the settings to use the `domain\fradmin1` account that was created in the steps above.
6. Restart the mount proxy client access daemon (CAD) service.

When you have set up an account with suitable privileges:

- **Windows** Enter your credentials in the Data Protection for VMware vSphere GUI configuration wizard or notebook to enable the environment for file restore operations.
- **Windows** A file owner accesses the remote virtual machine (that contains the files to be restored) with Windows domain user credentials. These credentials are entered in the file restore interface during login. Domain user credentials verify that the file owner has permission to log in to the remote virtual machine and restore files into the remote virtual machine. These credentials do not require any special permissions.
- **Windows** If a file owner uses a Windows domain user account that limits access to specific computers (instead of access to all computers within the domain), ensure that the mount proxy system is included in the list of computers that are accessible to this domain user account. Otherwise, the file owner is unable to log in to the file restore interface.

Tape media prerequisites

File restore from tape media is not supported. File restore from disk storage is the preferred method.

Logging in to restore files

You can log in to the IBM Storage Protect file restore interface to restore your files with minimal assistance from the administrator.

About this task

When you log in to this interface, you can locate and restore your files at your convenience.

Procedure

1. Access the file restore interface by opening a web browser and entering the URL that you received from your administrator.
2. Enter the network name or IP address of the virtual machine that contains your files. For example, `myhost.mycompany.com`.
3. Enter the user account that you use to access your files.

Windows

Use the `Windows_domain_name\user_name` format.

4. Enter the user account password and click **Log in**.

Restoring files from a virtual machine backup

Locate your files and restore them to a preferred location.

Before you begin

Ensure that you are logged in to the IBM Storage Protect file restore interface. A backup must exist before you can restore your files.

About this task

Only those files and directories for which you have permission to view on the operating system are visible.

Procedure

1. Select a backup by completing the following steps:
 - a) Click a date in the calendar.
 - b) If necessary, select a time in the **Available backups** field.
 - c) Click **Choose backup**.

The virtual machine disks or directories are displayed in the table.

2. Optional: If the default backup is not the one you want, select a different backup by completing the following steps:
 - a) Click the calendar.
 - b) Click a date in the calendar.
 - c) If necessary, select a time.
 - d) Click **Change backup**.

Restriction: If you change the backup date or time, any file selections that you made are lost. However, the new backup loads to the directory where you previously explored. If that directory is unavailable, the backup loads to the top directory.

The virtual machine disks or directories are displayed in the table.

3. To select files to restore, complete the following steps:
 - a) Click a disk or directory in the table to view the subdirectories and files.
 - b) Optional: To search for a file in the current directory and subdirectories, type a name in the **Search** field and press **Enter**. The results are displayed in the order they are found.

c) Select one or more files and directories to restore.

If you select a directory that has no contents, the empty directory is not restored.

4. Select where to restore files.

- To restore files and directories to the original location, select **Restore to > Original Location**.
- To restore files and directories to a different location, select **Restore to > Alternate Location**.


5. After you make your selections, click **Restore**.

If you are restoring files to an alternative directory, select an existing directory on your virtual machine or create a directory to place restored files. Then, click **Restore**.

If a file with the same name exists, the restored file's original modification date and time is added to the file name. Subsequent restores of the same file contain a number (*_N*) after the original modification date and time.

For example: t2.2015-03-07-07-28-03_1.txt

What to do next

Click the restore icon () to view information about active and recent restores. By default, information is kept for 7 days after a restore completes.

If a restore completed with an error or warning, view additional information by clicking **Details**. To save the error or warning information, click **Export** and save the information in .CSV format.

Chapter 4. Protection for in-guest applications

Data Protection for VMware protects Microsoft Exchange Server, Microsoft SQL Server, and Active Directory Domain Controllers that run inside VMware VM guests in a VMware vSphere environment.

Microsoft Exchange Server data protection in VMware environments

For Microsoft Exchange Server workloads that are running in a VMware ESXi virtual guest machine, you can take application-consistent backups of virtual machines that are hosting Microsoft Exchange Server data. You can recover database-level and mailbox-level backups from a virtual machine.

Before you back up data, identify your recovery point objective (RPO). The *RPO* helps you decide how often to back up data and affects the cost that is associated with data backups.

For example, you can schedule frequent virtual machine backups for necessary recovery points. The recovery point of a virtual machine backup is the time of the backup. While change-block tracking and data deduplication offer savings, virtual machine backups can be expensive if you create and delete many virtual machine snapshots.

Most traditional in-guest data protection methods provide appropriate RPOs, but these in-guest methods lose the efficiencies that are introduced by backing up data at the virtual machine level.

You can use Data Protection for VMware and Data Protection for Microsoft Exchange Server to back up data at a virtual machine level in a way that preserves backup efficiencies.

To protect Microsoft Exchange Server data in a VMware environment, ensure that the following products are installed and configured on your system:

- Data Protection for VMware 8.1.27
- Data Protection for Microsoft Exchange Server 8.1.11

These software offerings work together to protect Microsoft Exchange Server data in a VMware environment when no other software products are used to back up Microsoft Exchange Server data.

For permission required to back up and restore application data for Microsoft Exchange Server, see [technote 1647986](#).

Application protection is supported for VMware VMs in a VMware vSphere environment only.

Configuring the software for Exchange Server data protection in a VMware environment

To protect Microsoft Exchange Server workloads that are running in a VMware ESXi virtual guest machine, install and configure Data Protection for VMware. Then, install and configure Data Protection for Microsoft Exchange Server.

Before you begin

The configuration instructions that follow are based on a configuration scenario that might not match your environment. Adjust the configuration for your environment.

The following list summarizes the scenario for quick reference:

Windows host name

EXC10

VSS requestor node name

EXC10_VSS

Data Protection for Microsoft Exchange Server node name

EXC10_EXC

Virtual machine name

vm_exc10

Data mover node names

datamover10 and datamover20

Datacenter node name

datacenter10

VM file space

\VMFULL-vm_exc10

About this task

The following details describe the scenario that is used.

- A single Microsoft Exchange Server database on a virtual machine that is named vm_exc10 must be recovered.
- Virtual machine vm_exc10 is protected by Data Protection for VMware by using the node name datacenter10. This node name in the IBM Storage Protect server represents the vSphere datacenter. The data mover nodes are called datamover10 and datamover20.
- The virtual machine guest is configured with the virtual machine name of vm_exc10 and the Microsoft Windows host name is EXC10.
- Data Protection for Microsoft Exchange Server is installed in the virtual guest machine and is configured to the IBM Storage Protect server to use node name EXC10_EXC.
- The IBM Storage Protect client in the virtual guest machine is configured as the VSS requestor node and is using the node name EXC10_VSS.

Procedure

1. Follow the installation and configuration instructions that are provided with each software package.

If you install Data Protection for Microsoft Exchange Server before Data Protection for VMware, you cannot specify the VMware datacenter node in the Data Protection for Microsoft Exchange Server configuration wizard because the field is disabled.

Tip: In this case, reconfigure Data Protection for Microsoft Exchange Server to enable the VMware datacenter node after Data Protection for VMware is installed.

Data Protection will keep this VMware datacenter node value as the VEBACKUPNode option in the .cfg configuration file.

Do not forget to remove this option manually if you intend to uninstall Data Protection for VMware from the virtual machine hosting Data Protection for Exchange or SQL.

2. Complete the tasks in this checklist:
 - Verify that Microsoft Exchange Server databases and mailboxes are hosted on VMware virtual disks.
 - Verify that no Exchange Server database is hosted on raw device mapped (RDM) disks in physical compatibility mode, independent disks, or on disks that are attached directly to the guest through in-guest iSCSI.
3. Outside of the VM guest, in the data mover, configure Data Protection for VMware to protect Microsoft Exchange Server databases and mailboxes.
4. Within the guest, take the following actions:
 - Verify that the Data Protection for VMware recovery agent command-line interface is configured to work with the recovery agent on the guest virtual machine.
 - Configure Data Protection for Microsoft Exchange Server to complete backup and restore operations from a virtual machine backup.

Related tasks

[“Configuring Data Protection for VMware” on page 59](#)

You must configure Data Protection for VMware to preserve Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft Exchange Server data.

[“Configuring Data Protection for Microsoft Exchange Server” on page 61](#)

After you configure Data Protection for VMware and verify that you created a virtual machine backup that is suitable for recovery of a single Microsoft Exchange Server database, configure Data Protection for Microsoft Exchange Server in the guest virtual machine.

Configuring Data Protection for VMware

You must configure Data Protection for VMware to preserve Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft Exchange Server data.

About this task

Data Protection for VMware provides application consistency when you back up virtual machines that are hosting Microsoft Exchange Servers. With these backups, you can recover the virtual machine with Microsoft Exchange Server in a consistent state.

To recover only selected databases or mailboxes from this type of backup with IBM Storage Protect, without having to recover the entire virtual machine, preserve information about the state of the Microsoft Exchange Server at the time of the virtual machine snapshot and backup. This information is collected as part of the Microsoft Volume Shadow Copy Services (VSS) interaction that occurs during a virtual machine snapshot.

For Data Protection for VMware to collect the Microsoft VSS metadata for Microsoft Exchange Server, you must configure Data Protection for VMware to obtain this information from the virtual machine during the backup operation.

Procedure

1. Configure Data Protection for VMware to preserve the Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft Exchange Server data.
 - a) Locate the options file for the Data Protection for VMware data mover. On Windows systems, the options file is `dsm.opt`. On Linux systems, the options file is `dsm.sys`.
 - b) Specify the `INCLUDE.VMTSMVSS` option for the virtual machine.

You must set this option for virtual machine backups to preserve the Microsoft VSS metadata information. For example, specify `INCLUDE.VMTSMVSS vm_display_name` where `vm_display_name` refers to the name of the virtual machine as shown in the VMware vSphere Client and vSphere Web Client.
 - c) Optional: Back up a passive copy of a database that is part of an Exchange Server Database Availability Group (DAG). Specify the `vmpreferdagpassive yes` option for the virtual machine.

Backing up the passive copy typically reduces the performance impact to the active copy in the production database. If no valid passive copy is available, the active database copy is backed up.
 - d) Verify that the Virtual Machine Disks (VMDK) that host the Microsoft Exchange Server database are not being excluded from the virtual machine backup operation.

Repeat the preceding steps for all data movers that protect virtual machines that are hosting Microsoft Exchange Server.
2. On each data mover, for example, `datamover10`, store the guest virtual machine credentials to Data Protection for VMware by running the following command from the IBM Storage Protect backup-archive client command line:

```
dsmc set password -type=vmguest vm_guest_display_name
guest_admin_ID guest_admin_pw
```

This command stores the guest virtual machine credentials, which are encrypted on the system that hosts the data mover. The following minimum permissions are required for *guest_admin_ID* *guest_admin_pw*:

- Backup rights: Microsoft Exchange Server 2010 and 2013: Organization Management permissions (membership in the management role group, Organization Management).

What to do next

You can verify the virtual machine backup configuration and ensure that the VMDKs are included. In addition, you can view other parameters by issuing the **backup** command with the preview option from the data mover, for example:

```
dsmc backup vm vm_display_name -preview -asnode=datacenter_node
```

You can also use the IBM Storage Protect scheduler to schedule periodic backups of your virtual machines. You can also back up the virtual machine that is hosting the Microsoft Exchange Server by using the data mover command line:

```
dsmc backup vm vm_display_name -asnode=datacenter_node
```

Verifying that the configuration backs up data that can be restored

Before you can restore individual Microsoft Exchange Server databases from a Data Protection for VMware virtual machine backup, you must complete at least one successful virtual machine backup. For the restore operation to work, the backup must contain Microsoft Exchange Server database metadata.

Procedure

1. Issue the following data mover **query** command on one of the data mover nodes:

```
dsmc query vm vmname -detail -asnode=datacenter_node
```

where:

- *vmname* specifies the name of the virtual machine
- *datacenter_node* specifies the name of the datacenter node

.

2. In the command output, look for the following details:

```
Application(s) protected: MS Exchange 2013 (database-level recovery)
```

Ensure that **Excluded** is not indicated in any Virtual Machine Disk (VMDK) status fields for virtual disks that host Microsoft Exchange Server database files. The **Excluded** status indicates that one or more of the VMDKs that are required to recover a Microsoft Exchange Server database are not being protected. For example:

```

Query Virtual Machine for Full VM backup
# Backup Date      Mgmt Class  Size      Type      A/I      Virtual Machine
-----
1 02/20/2015      STANDARD    43.94GB   IFFULL    A        vm_exc10
  12:43:59

    Size of this incremental backup: n/a
    Number of incremental backups since last full: 0
Amount of extra data: 0
Object fragmentation: 0
Backup is represented by: 328 objects
Application protection type: TSM VSS
Application(s) protected: MS EXC 2013 (database-level recovery)
VMDK[1]Label: Hard Disk 1
VMDK[1]Name: [ess800_dev2] vm_exc10/vm_exc10 .vmdk
VMDK[1]Status: Protected
...
VMDK[6]Label: Hard Disk 6
VMDK[6]Name: [ess800_dev2] vm_exc10/vm_exc10_5.vmdk
VMDK[6]Status: Protected

```

Configuring Data Protection for Microsoft Exchange Server

After you configure Data Protection for VMware and verify that you created a virtual machine backup that is suitable for recovery of a single Microsoft Exchange Server database, configure Data Protection for Microsoft Exchange Server in the guest virtual machine.

Procedure

1. Log on to the virtual machine that hosts the Microsoft Exchange Server database.
2. Verify that the following packages are installed:
 - The IBM Storage Protect recovery agent, recovery agent command-line interface (CLI), and license (from the Data Protection for VMware product package)
 - IBM Storage Protect data mover
 - Data Protection for Microsoft Exchange Server

You can install the recovery agent, CLI, license, and data mover together by using the Data Protection for VMware installation program. To install the packages together, select the following advanced installation option: **Install a complete data mover for in-guest application protection**. Data Protection for Microsoft Exchange Server is installed separately.
3. Configure Data Protection for Microsoft Exchange Server by using the IBM Storage Protect configuration wizard. When you open the **IBM Storage Protect Node Names** page of the wizard, enter the VMware datacenter node name, Microsoft Exchange Server node name, and VSS requestor node name. If the datacenter node name field is disabled, the recovery agent is not installed correctly.
4. After Data Protection for Microsoft Exchange Server is configured, verify that the **Configuring Recovery Agent** rule status indicates Passed.
5. Log on to a data mover instance and complete the following steps.

Do not repeat these steps for all data mover instances.

 - a) Copy the contents of the data mover options file `dsm.opt` to a temporary file named `dsm.setaccess.opt` and make the following changes to the file. Do not make these changes in the `dsm.opt` file.
 - i) Delete any line that contains an **ASNODE** entry.
 - ii) Set the **NODENAME** option to the VMware datacenter node name. For example:

```
NODENAME datacenter10
```

Tip: If the `dsm.opt` file does not contain **ASNODE** entries and the **NODENAME** option is set to the correct data center node, you can use this file rather than creating the `dsm.setaccess.opt` file.

- b) From the datacenter node that was defined in the **NODENAME** option, issue the **set access** command to grant the VSS requestor node access to the virtual machine backups as shown in the following example.

You must complete this step because the VSS requestor node accesses the virtual machine backups on behalf of Data Protection for Microsoft Exchange Server.

If the password for the datacenter node is unknown when you run the **set access** command, you will receive an error message and the IBM Storage Protect server administrator must reset the password to issue the command.

Example

The following example shows the required parameters for the **set access** command. The parameters specify the virtual machine name (vm_exc10), the VSS requestor node name (EXC10_VSS), and the name of the options file that defines the datacenter node name (dsm.setaccess.opt).

```
dsmc set access backup -type=VM vm_exc10 EXC10_VSS -optfile=dsm.setaccess.opt
ANS1148I "Set Access" command successfully completed.
```

The following example shows the results of the **query access** command, which shows the backup access authorization for the VSS requestor node.

```
dsmc query access
Node name: datacenter10
Type      Node      User      Path
-----
Backup    EXC10_VSS  *         \VMFULL-vm_exc10\*\*
```

ANS1148I "Query Access" command completed successfully

Managing backups

After you configure Data Protection for Microsoft Exchange Server, you can schedule virtual machine backups and separately, you can update the mailbox information in Exchange Server database backups on the virtual machine.

Scheduling virtual machine backups

To ensure that your data is protected, schedule virtual machine backups.

Before you begin

Before you back up virtual machines that are hosting Microsoft Exchange Server databases, mount the databases.

By default, the maximum size allowed for a VMDK in a backup operation is 2 TB. However, the maximum is 8 TB. To increase the maximum size, use the `vmmaxvirtualdisks` option. For more information, see [Vmmaxvirtualdisks](#).

About this task

During backup processing, Data Protection for VMware bypasses a guest Microsoft Exchange Server database that is dismounted, corrupted, suspended, or not in a healthy state in a Database Availability Group (DAG). Databases in such invalid states are excluded from virtual machine backups and are not available to restore.

Procedure

1. Log on to the Data Protection for VMware user interface.
2. Click the **Backup** tab.

3. Click **Create Schedule** to specify a backup schedule name, source (the virtual machines to include in the backup schedule), and other scheduling options.
4. Verify that the source of the schedule includes the virtual machines that are hosting Microsoft Exchange Server.
5. Verify that one of the following services is running:
 - If you are using scheduler that is managed by a Client Acceptor Daemon (CAD), ensure that the CAD service is running on the data mover.
 - If you are using the stand-alone scheduler, ensure that the scheduler service is running.

Updating mailbox information in Microsoft Exchange Server backups

When you back up a virtual machine that is hosting Microsoft Exchange Server data, mailbox history is automatically uploaded with the virtual machine backup if Data Protection for Microsoft Exchange Server is detected on virtual machine.

About this task

Unless Data Protection for Microsoft Exchange Server is installed on the virtual machine, mailbox history information is not automatically updated in Exchange Server database backup operations. Automatic uploading of mailbox history might also be disabled by specifying the `VMBACKUPMAILBOXHISTORY No` in the data mover options file, that is, `dsm.opt` on Windows systems or `dsm.sys` on Linux systems.

You can manually update mailbox history information by using the Data Protection for Microsoft Exchange Server command-line interface.

Tip: Complete this task before you back up the virtual machines that contain Microsoft Exchange servers. In this way, you can ensure that you have consistent location information for the mailbox history and the mailboxes in database backups.

Procedure

1. To update only the mailbox history information in Exchange Server database backups, issue the **backup /UpdateMailboxInfoOnly** command as shown in the following example:

```
tdpexcc backup DB1 full /UpdateMailboxInfoOnly
```

where `DB1` is the database name, and `full` is the type of database backup.

Tip: To update information for all the mailboxes in the Exchange organization, specify an asterisk (*) character as the database name.

2. Optional: Verify that the mailbox information is updated correctly by completing the following steps.
 - a) Review the mailbox information for database backups on IBM Storage Protect server by issuing the **query /SHOWMAILBOXInfo** command as shown in the following example:

```
tdpexcc query tsm /showmailboxinfo
```

- b) Start Microsoft Management Console (MMC), and in the **Mailbox Restore** or **Mailbox Restore Browser** view, verify the list of updated mailboxes that are available to restore.

Verifying backups

After you create a backup, verify that you can query the virtual machine backups and the database backups from the Data Protection for Microsoft Exchange Server interface.

About this task

You can recover one or more Microsoft Exchange databases based on your recovery point objectives (RPO).

Procedure

1. From Microsoft Management (MMC), select a Microsoft Exchange Server.
2. Click the **Recover** tab.
3. Select **View > Databases**. A list of Microsoft Exchange Server database backups that can be restored is displayed.

Microsoft Exchange Server databases that are backed up with Data Protection for VMware are identified with the vmvss backup method.

Troubleshooting VSS backup operations on guest virtual machines

If you encounter a problem during Volume Shadow Copy Service (VSS) backup processing on a guest VM, try to reproduce the problem in your environment.

About this task

Contact IBM Support for further assistance if you have a problem that you are unable to solve by reproducing the issue or reviewing the information that follows.

VSS writer service causes a VM backup to fail

You can bypass any VSS writer that is causing a VM backup to fail and exclude it from the backup.

About this task

Before a VM backup, the VSS writer is in a stable state and has no errors. During VM backup processing, a VSS writer might encounter an error that causes the entire VM backup to fail.

For example, if the Microsoft Forefront Protection VSS Writer is installed on a guest VM, the VM backup fails and the VSS writer status changes to `Retryable error`, `Waiting for completion`, or a status other than `Stable`. Complete the following steps to exclude the writer service from the VM backup.

Procedure

1. In the VSS administrative command-line tool on the guest VM, list the VSS writers by issuing the **vssadmin list writers** command.
In the following command example, the Microsoft Forefront Protection VSS Writerservice is identified by writer name, ID, and instance ID:

```
Writer name: 'FSCVSSWriter'  
  Writer Id: {68124191-7787-401a-8afa-12d9d7ccc6ee}  
  Writer Instance Id: {f4cc5385-39a5-463b-8ab4-aafb2b35e21e}  
  State: [1] Stable  
  Last error: No error
```

2. In the datamover options file, `dsm.opt` or `dsm.sys`, add the `EXCLUDE.VMSYSTEMSERVICE` option followed by the *Writer Name* as shown in the following example.

```
EXCLUDE.VMSYSTEMSERVICE FSCVSSWriter
```

Tip: If the data mover machine is on a UNIX system, the option file is `dsm.sys`. If the guest VM and datamover machine use different language sets, specify the *Writer ID* or *Writer Instance Id* instead of the *Writer Name*.

For example:

```
EXCLUDE.VMSYSTEMSERVICE {68124191-7787-401a-8afa-12d9d7ccc6ee}
```

Results

The VM backup completes successfully even if the Microsoft Forefront Protection VSS Writer service is running on the guest VM.

No application protection file APPPROTECTIONDBINFO.XML and no warning messages for skipped databases

Under certain conditions, a dismounted Exchange 2010 Server database is skipped during a backup operation and no warning is issued.

About this task

When the following conditions exist during a VM backup of a guest VM with Exchange 2010 Server:

- The Exchange 2010 Server is not a member of a Database Availability Group (DAG).
- All Exchange 2010 Server databases are dismounted.

The following warning message is generated:

```
ANS4063W IBM Storage Protect application protection cannot copy
the application metafile 'APPPROTECTIONDBINFO.XML ' from the following VM: '<name_name>'.
Individual database restore from this backup is not supported.
```

```
ANS4063W IBM Storage Protect application protection cannot copy the
application metafile '-----L' from the following VM: '<vm_name>'.
Individual database restore from this backup is not supported.
```

In this situation, the VM backup is available for only full VM restore. Individual database restore from this VM backup is not available.

To prevent this situation, mount the Exchange 2010 Server databases before you start the VM backup operation.

When Exchange 2010 Server DAG databases or Exchange Server 2013 databases are dismounted, a VM backup operation of a guest VM generates the following warning message:

```
ANS2234W Restore from virtual machine backup is not available for
dismounted database <database>
```

For a dismounted Exchange 2010 Server database that is not a member of a DAG, IBM Storage Protect does not detect that the databases are dismounted. As a result, warning message ANS4063W is generated instead of ANS2234W.

Transaction error due to mixing of deduplicated and non-deduplicated data in the same transaction

Under certain conditions, a transaction error occurs when deduplicated and non-deduplicated data is mixed in the same transaction.

About this task

When data deduplication is enabled, a Data Protection for VMware backup with application protection of a virtual machine might generate the following error in the `dsmerror.log` file:

```
ANS0246E Issue dsmEndTxn and then begin a new transaction session.
ANS5250E An unexpected error was encountered.
  IBM Storage Protect function name : vmSendViaFile()
  IBM Storage Protect function      : Failed sending file
                                     /tmp/tsmvmbackup/fullvm/vmtsmvss/member1/IIS CONFIG WRITER.XML
  IBM Storage Protect return code   : 2070
  IBM Storage Protect file          : vmmigration.cpp (1383)
```

This error is recoverable and can be ignored. The error occurs when Data Protection for VMware attempts to send the XML file (that was excluded from deduplication due to its small size) in the same transaction with deduplicated data. Data Protection for VMware resends the XML file (identified in the error message) in a new transaction.

Verifying that virtual machine backups do not exclude Microsoft Exchange Server volumes

The volumes in Virtual Machine Disks (VMDKs) must contain the Microsoft Exchange Server databases that are not excluded from the Data Protection for VMware backup processing.

About this task

The databases cannot be on physical compatibility mode raw device mapping (RDM) disks, independent disks, or on disks that are attached directly to the guest operating system through iSCSI.

Procedure

1. Ensure that any EXCLUDE.VMDISK statements in the Data Protection for VMware data mover that is used to back up the virtual machine do not inadvertently exclude VMDKs that are hosting volumes that contain Microsoft Exchange Server files, file space, database, and mailboxes.

For example:

- vm_exc10.vmdk contains logical volume C:
 - vm_exc10.vmdk contains logical volumes E: and F:
 - The label for vm_exc10_1.vmdk is *Hard Disk 1*.
 - The label for vm_exc10_2.vmdk is *Hard Disk 2*.
 - The Microsoft Exchange Server database files to be backed up are on the E: and F: drive.
2. Verify that no statements exclude vm_exc10_2.vmdk from the virtual machine backup by ensuring that the data mover does not contain the following or similar statements:

```
EXCLUDE.VMDISK VM_EXC10 "Hard Disk 2"  
EXCLUDE.VMDISK * "Hard Disk 2"
```

Alternatively, if you exclude most hard disks, you must explicitly include the virtual machine disks by using one of the following statements:

```
INCLUDE.VMDISK VM_EXC10 "Hard Disk 2"  
INCLUDE.VMDISK * "Hard Disk 2"
```

Include and exclude statements are processed from bottom to top as they are displayed in the dsm.opt file. To achieve the goal, enter the statements in the correct order.

You can specify the exclusion and inclusion of a virtual machine disk from the command-line interface:

```
dsmc backup vm "VM_EXC10:-vmdisk=Hard Disk 2" -asnode=datacenter10
```

Restoring data

After you back up data, you can recover the data based on a recovery point objective (RPO).

A recovery operation restores a full backup of the Microsoft Exchange Server database or mailbox from the Data Protection for VMware backup.

If you restore the entire virtual machine, all Microsoft Exchange Server databases and mailboxes on the virtual machine are restored and recovered to the point of the virtual machine backup.

Starting the Microsoft iSCSI Initiator Service

The iSCSI protocol is used to mount the disks that are used for a recovery operation. Ensure that the Microsoft iSCSI Initiator Service is started and is set to the automatic startup type on the system where the data is to be restored.

Procedure

1. In the Windows **Services** list, right-click **Microsoft iSCSI Initiator Service**.
2. Click **Properties**.
3. On the **General** tab, set the following options:
 - a) In the **Startup type** list, select **Automatic**.
 - b) Click **Start**, and then click **OK**.

Results

In the **Services** list, **Microsoft iSCSI Initiator Service** shows a status of **Started** and the startup type is **Automatic**.

Restoring database backups by using the graphical user interface

You can recover a full Microsoft Exchange Server database backup from a virtual machine backup by using the Data Protection for Microsoft Exchange Server graphical user interface.

Procedure

1. To start a full database recovery from a virtual machine, start Microsoft Management (MMC). In the navigation pane, expand the **Protect and Recover** node and select a Microsoft Exchange Server server.
2. On the **Recover** tab, select **Database Restore**. All backups, including all database backups from a virtual machine backup, are listed.
3. Select a full database backup to restore.
4. In the **Actions** pane, click **Restore**.

Restoring backups of another virtual machine

By using Data Protection for Microsoft Exchange Server, you can access backups of another virtual machine on IBM Storage Protect server and restore the backup.

About this task

You can restore database and mailbox backups to a different Database Availability Group (DAG) node than the original backup node.

The following scenario assumes that you have Exchange virtual machines in your virtual environment: vm1 and vm2. You want to enable Data Protection for Microsoft Exchange Server on vm2 to access and restore database and mailbox backups on vm1 and vm2.

Procedure

1. Configure self-contained application protection to protect Microsoft Exchange Server data on vm1 and vm2.
2. Back up vm1 and vm2 by issuing the **dsmc backup vm** command on the command-line interface.
3. On vm2, install Data Protection for Microsoft Exchange Server and configure the software for Exchange Server data protection in a VMware environment.
4. To enable Data Protection for Microsoft Exchange Server on vm2 to access backups on vm1 and vm2, issue the **set access** command as shown in the following examples:

```
dsmc set access backup -type=vm vm1 vm2_vss
```

```
dsmc set access backup -type=vm vm2 vm2_vss
```

5. Restore database or mailbox backups on vm1 or vm2.

Related tasks

[“Configuring Data Protection for VMware” on page 59](#)

You must configure Data Protection for VMware to preserve Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft Exchange Server data.

[“Configuring the software for Exchange Server data protection in a VMware environment ” on page 57](#)

To protect Microsoft Exchange Server workloads that are running in a VMware ESXi virtual guest machine, install and configure Data Protection for VMware. Then, install and configure Data Protection for Microsoft Exchange Server.

[“Configuring Data Protection for Microsoft Exchange Server” on page 61](#)

After you configure Data Protection for VMware and verify that you created a virtual machine backup that is suitable for recovery of a single Microsoft Exchange Server database, configure Data Protection for Microsoft Exchange Server in the guest virtual machine.

Restoring mailbox data

Data Protection for Microsoft Exchange Server backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. For more information, see the topic that explains security requirements for backup and restore operations in the *Data Protection for Microsoft Exchange Server Installation and User's guide*.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) server, ensure that you configure the SMTP server before you start a restore operation. To set the configuration in the Management Console, right-click **Dashboard** in the tree view and select **Properties**. From the **E-mail** property page, enter the SMTP server and port.

About this task

- In Exchange Server 2013, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
 - To restore an Exchange 2013 public folder mailbox, the Exchange user must have the **Public Folders** management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.

- You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore

operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.

- In Exchange Server 2010 or later, you can restore an archive mailbox or a part of the mailbox, for example, a specific folder. You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

If you enable a user mailbox to be archived, ensure that the user is logged on to that mailbox at least once before you complete a backup and restore operation on the mailbox.

- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.
- By default, Data Protection for Microsoft Exchange Server restores the latest backup that is available for the specified mailbox.

The amount of time that it takes to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

Procedure

1. Start Microsoft Management Console (MMC) and select **Exchange Server** in the navigation tree.
2. On the **Recover** tab for the Exchange Server instance, select the **Mailbox Restore** view.
3. Select one or more mailboxes to restore. A list of mailboxes that are backed up is displayed.

If you restore mail to a Unicode personal folder (.pst) file, or you restore a mailbox that is deleted or re-created after the time of the backup, Data Protection for Microsoft Exchange Server requires a temporary mailbox to store the mailbox messages. Create a temporary mailbox by setting the **Alias of temporary mailbox** option on the **Properties** page, under the **General** tab.



Attention: Ensure that the temporary mailbox that you create is on a database with enough disk storage capacity to accommodate all of the mailbox items that you are restoring.

4. Optional: Optional: To restore individual messages instead of the entire mailbox, select **Item-Level Mailbox Filters**.
 - a) Click **Show Filter Options** and **Add Row**.
 - b) In the **Column Name** field, click the down arrow and select an item to filter.
 - You can filter public mailbox folders only by the **Folder Name** column.
 - You can filter Unicode .pst files only by **Backup Date**, **Folder Name** and **All Content** filters.
 - You can filter by backup date, and click the default date and time to edit the table cell. To change the date, click the arrow at the end of the cell. The calendar date selection tool is displayed. After you select a date, to display the date in the field, press **Enter**. To edit the time, enter the time by using the 12-hour clock time convention such as 2 p.m.

When you specify a backup date, Data Protection for Microsoft Exchange Server searches for a backup that corresponds to that exact date. If a backup with that exact date is not found, Data Protection for Microsoft Exchange Server selects the first backup after that date.
 - c) In the **Operator** field, select an operator.
 - d) In the **Value** field, specify a filter value.
 - e) If you want to filter on more items, click **Add Row**.
5. Specify the restore options by clicking **Show Restore Options**.

<i>Table 3. Database restore options</i>	
Task	Action
Keep Recovery Database After Restore	Use this option to retain a recovery database after a mailbox restore operation is complete. The default value is No . If you set the value to Yes , Data Protection for Microsoft Exchange Server automatically retains the recovery database after mailbox restore processing.
Mailbox	If the alias of the mailbox to restore is not displayed in the list of mailboxes, specify the alias. This option overrides any selected mailboxes.
Mailbox Original Location	Use this option only if the mailbox was deleted or re-created since the time of the selected backup, and mailbox history is disabled. Specify the Exchange Server and the database where the mailbox was at the time of the backup. Use the following format: server-name, db-name, for example, serv1, db1.
Mark Restored Messages As Unread	Use this option to automatically mark the mailbox messages as unread after the restore operation is completed. The default value is Yes .
Use Existing Recovery Database	Use this option to restore the mailbox from an existing recovery database. The default value is Yes . If you set the value to No and a recovery database is mounted on the server before you restore the mailbox, Data Protection for Exchange Server automatically removes the recovery database during mailbox restore processing.

6. To complete the restore operation, click one of the following **Restore** options.

<i>Table 4. Restore options</i>	
Task	Action
Restore Mail to Original Location	Select this action to restore mail items to their location at the time of the backup operation.
Restore Mail to Alternate Location	Select this action to restore the mail items to a different mailbox.

Table 4. Restore options (continued)	
Task	Action
Restore Mail to non-Unicode PST file	<p>Select this action to restore the mail items to a non-Unicode personal folders (.pst) file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory.</p> <p>If the .pst file exists, the file is used. Otherwise, the file is created.</p> <p>Restriction: The contents of each folder cannot exceed 16,383 mail items.</p>
Restore Mail to Unicode PST file	<p>Select this action to restore the mail items to a Unicode .pst file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location.</p> <p>You can enter a standard path name (for example, c:\PST\mailbox.pst) or a Universal Naming Convention (UNC) path (for example, \\server\c\$\PST\mailbox.pst). When you enter a standard path, the path is converted to a UNC path. If the UNC path is a non-default UNC path, enter the UNC path directly.</p> <p>Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created.</p>

Table 4. Restore options (continued)	
Task	Action
Restore Public Folder Mailbox	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>
Restore Mail to Archive Mailbox	<p>This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternate archive mailbox.</p> <p>You can filter the archive mailbox and restore a specific mailbox folder. In the Folder to be restored field, enter the name of the folder in the archive mailbox that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>In the Target archive mailbox field, specify the archive mailbox destination that you want to restore to.</p>

Tip: Because a status indicator does not appear in MMC during the restore operation, you might assume that the operation has stopped because it is taking a long time to complete. However, depending on the amount of data, a restore operation can take several hours.

Restoring relocated and deleted mailboxes

The backup solution for restoring mailboxes that are relocated and deleted after a virtual machine backup consists of Data Protection for VMware and Data Protection for Microsoft Exchange Server.

Before you begin

Decide where the mailbox data from the deleted mailbox is to be restored.

If you restore mail to a Unicode personal folder (.pst) file, or you restore a mailbox that is deleted or re-created after the time of the backup, Data Protection for Exchange Server requires a temporary mailbox to store the mailbox messages. Create a temporary mailbox by setting the **Alias of temporary mailbox** option on the **Properties** page, under the **General** tab.



Attention: Ensure that the temporary mailbox that you create is on a database with enough disk storage capacity to accommodate all of the mailbox items that you are restoring.

About this task

When you restore the backups, and complete a full database restore operation from the backup, Data Protection for VMware restores the files to their original location.

If database or log files are relocated during the backup cycle, Data Protection for Microsoft Exchange Server restores the files in their original locations.

If any databases or mailboxes were created during the backup cycle, Data Protection for Microsoft Exchange Server re-creates the new files. If database or log files were deleted during the backup cycle, those files are not restored.

Procedure

Complete one of the following actions:

- Restore the deleted mailbox data to the original location. Before you run the mailbox restore operation, re-create the mailbox that is using Exchange.

If the backup that contains the deleted mailbox was created with a version of Data Protection for Microsoft Exchange Server earlier than version 6.1, or if the mailbox history is disabled, and the mailbox was relocated after the time it was backed up, you must specify the Exchange Server and the database where the mailbox was at the time of backup. Use the **Mailbox Original Location** option in the GUI to specify this information. Alternatively, issue the **restoremailbox** command parameter, /**MAILBOXORIGLOCATION**.

- Restore the deleted mailbox data into an active alternative mailbox in an online Exchange Server.
- Restore the deleted mailbox data into an Exchange Server personal folders (.pst) file.

Restoring mailbox messages interactively with the Mailbox Restore Browser

You can use the **Mailbox Restore Browser** to interactively restore a mailbox or items from a mailbox on Exchange Server.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) Server, ensure that you configure the SMTP Server before you attempt a restore operation. Set the configuration in Microsoft Management Console (MMC) by right-clicking **Dashboard** in the tree view and selecting **Properties**. Then, in the **E-mail** property page, enter the SMTP server and port.

About this task

- In Exchange Server 2013, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder. However, you cannot restore individual messages in a public folder by using the Mailbox Restore Browser interface.
 - To restore an Exchange 2013 public folder mailbox, the Exchange user must have the **Public Folders** management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.
- If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
- You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.
 - By default, Data Protection for Exchange Server restores the latest backup that is available for the specified mailbox.

Restriction: Only mailboxes within the same database can be restored in a single mailbox restore action.

Procedure

1. Start MMC.
2. Under the **Protect and Recover Data** node in the navigation tree, select **Exchange Server**.
3. On the **Recover** panel, click **View > Mailbox Restore Browser**.
4. In the **Select Source** window, specify the mailbox that you want to restore.

Choose from the actions in the following table:

Table 5. Selecting mailboxes to restore	
Task	Action
Browse mailboxes and select one to restore	<ul style="list-style-type: none">a. From the drop-down list, select Browse Mailboxes.b. Select a mailbox.c. Click OK. <p>Tip: Use the Search field to filter the mailboxes. You can also sort the mailboxes by columns.</p>
Specify a mailbox to restore by name	<ul style="list-style-type: none">a. In the Mailbox Name field, enter the name of the mailbox to restore.b. Click OK.

<i>Table 5. Selecting mailboxes to restore (continued)</i>	
Task	Action
Restore a mailbox backup that was created at a specific time	<ul style="list-style-type: none"> a. In the Backup Date/Time field, click the default date and time to edit the details. b. To change the date, click the calendar icon, select a date, and press Enter. c. To change the time of day, use the 12-hour system convention such as 2 p.m. d. Click OK.
Review the mailbox backups that are available to restore before you complete the restore operation	<ul style="list-style-type: none"> a. From the drop-down list, select Browse Mailboxes. b. Select a mailbox for which backups exist. c. From the Available Database Backups list, review the backups that are available for the mailbox and select a backup version to restore. d. Ensure that the Backup Date/Time field reflects the time stamp for the selected mailbox backup. e. Click OK.
Restore a mailbox that was deleted or re-created after the time of the backup	<p>In the Actions pane, click Properties, and on the General page, enter the temporary mailbox alias.</p> <p>Tip: If you do not enter the alias, the mailbox restore operation uses the administrator mailbox as a temporary storage location.</p>
Browse all databases in a backup	<ul style="list-style-type: none"> a. From the drop-down list, select Browse Databases. b. Select a database. c. Click OK. <p>Tip: Use the Search field to filter the databases. You can also sort the mailboxes by columns.</p>

After the selected mailbox is restored to the recovery database, the restored mailbox and folders are displayed in the results pane.

5. In the results pane, browse the folders and messages that are contained within the selected mailbox. Choose from the following actions to select the mailbox, folder, or message to restore:

<i>Table 6. Previewing and filtering mailbox items</i>	
Task	Action
Preview mailbox items	<ul style="list-style-type: none"> a. Select a mailbox item to display its contents in the preview pane. b. When an item contains an attachment, click the attachment icon to preview its contents. Click Open or save the item by clicking Save.

Table 6. Previewing and filtering mailbox items (continued)

Task	Action
Filter mailbox items	<p>Use the filter options to narrow the list of folders and messages in the result pane.</p> <ol style="list-style-type: none"> Click Show Filter Options and Add Row. Click the down arrow in the Column Name field and select an item to filter. You can filter by folder name, subject text, and so on. <p>You can filter public mailbox folders only by the Folder Name column.</p> <p>When you select All Content, the mailbox items are filtered by attachment name, sender, subject, and message body.</p> <ol style="list-style-type: none"> In the Operator field, select an operator. In the Value field, specify a filter value. If you want to filter on more items, click Add Row. Click Apply Filter to filter the messages and folders.

6. In the **Actions** pane, click the folder or messages restore task that you want to run.

If you click **Save Mail Message Content**, which becomes available only when a message is selected in the preview pane, a Windows Save File window is displayed. Specify the location and message name and click **Save**.

The **Restore Progress** window opens and shows the progress of the restore operation. IBM Storage Protect Snapshot restores the mailbox backup to its original mailbox location.

7. To restore a mailbox or mailbox item to either of the following locations, complete the following steps. Choose from the actions in the following table:

Table 7. Restoring a mailbox to another mailbox or .pst file

Task	Action
Restore a mailbox or mailbox item to a different mailbox	<ol style="list-style-type: none"> On the Actions pane, click Open Exchange Mailbox. Enter the alias of the mailbox to identify it as the restore destination. Drag the source mailbox to the destination mailbox on the results pane.
Restore a mailbox to an Outlook personal folders (.pst) file	<ol style="list-style-type: none"> On the Actions pane, click Open PST File. When the Windows File window opens, select an existing .pst file or create a .pst file. Drag the source mailbox to the destination .pst file on the results pane. <p>Restriction: You can use the Mailbox Restore Browser only with non-Unicode .pst files.</p>

Table 7. Restoring a mailbox to another mailbox or .pst file (continued)	
Task	Action
Restore Public Folder Mailbox	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>

In the Actions pane, the **Close Exchange Mailbox** and **Close PST File** tasks are displayed only when a destination mailbox or .pst file is opened.

- Optional: Remove the recovery database by clicking **Close Mailbox to Restore**.

This option is displayed only after a recovery database is created. Data Protection for Exchange Server removes the recovery database and cleans up the restored files. If you do not select **Close Mailbox to Restore**, the recovery database is not removed even if you exit MMC.

If MMC also detects a recovery database that is created outside of Data Protection for Exchange Server, it automatically connects to it. When you complete your mailbox restore tasks, you must manually remove the recovery database. You cannot use the **Close Mailbox to Restore** option.

Restoring data by using the command-line interface

If you prefer, you can use the command-line interface to start a full Microsoft Exchange Server database recovery from a virtual machine.

Procedure

- Issue the **query** command to find the database full backups.

The following example finds all backups for the Microsoft Exchange Server database called exc_db10.

```
tdpexcc q tsm exc_db10 IBM Storage Protect for Mail:
Data Protection for Microsoft Exchange Server Version 8, Release 1, Level 0.0
...
Querying IBM Storage Protect server for a list of
data backups, please wait...

Connecting to IBM Storage Protect Server as node "exc_db10"...
Connecting to Local DSM Agent "exc"...
Using backup node "exc_db10"...

Exchange Server      : exc
Database              : exc_db10
```

Backup Date	Size	S	Type	Loc	Object Name
10/15/2014 19:17:26	5.40 B	A	full	Srv	20141015191726 (VMVSS)

The operation completed successfully. (rc = 0)

2. To restore the database without applying transaction logs, issue the database **restore** command as shown in the following example:

```
TDPEXCC RESTore databaseName FULL /BACKUPDestination=TSM
/BACKUPMethod=VMVSS
```

The following sample output results when you issue the command with the Microsoft Exchange Server database called `exc_db10`.

```
TDPEXCC RESTore exc_db10 FULL /BACKUPDestination=TSM /BACKUPMethod=VMVSS
IBM Storage Protect for Mail:
Data Protection for Microsoft Exchange Server
Version 8, Release 1, Level 0.0 (C) Copyright
IBM Corporation 1997, 2016. All rights reserved.

Connecting to IBM Storage Protect Server as node "exc_db10"...

Connecting to Local DSM Agent "exc"...
Using backup node "exc_db10"...

Starting Microsoft Exchange restore...
Beginning VSS restore of "exc_db10"...

Restoring "exc_db10" via file-level copy from snapshot(s).
This operation could take a while, please wait

...

The operation completed successfully. (rc = 0)
```

You can restore the database to a different location by adding the **/INTODB** parameter. For example:

```
TDPEXCC RESTore TestDB1 FULL /INTODB=Test2
/BACKUPDestination=TSM /BACKUPMethod=VMVSS
```

What to do next

You can restore inactive backups by using the Data Protection for Microsoft Exchange Server command-line interface, **TDPEXCC**. When you issue the **restore** command, specify the database object name for the specific backup.

To obtain the database object name, issue the following command:

```
tdpexcc q tsm dbname full /all
```

After you have the database object name value, specify the database object name on the **/OBJect=objectname** parameter of the **TDPEXCC restore** command, where *objectname* is the database object name. For example:

```
TDPEXCC RESTore db44 FULL /OBJect=20140311131051 /BACKUPDestination=TSM
/BACKUPMethod=VMVSS
```

Restoring data by using Windows PowerShell cmdlets

If you prefer, you can use Windows PowerShell cmdlets with IBM Storage Protect Snapshot to start a full Microsoft Exchange Server database recovery from a virtual machine.

Procedure

1. Issue the query cmdlet to find the database full backups.
For example, to find all of the database full backups, enter the following command:


```
Get-DpExcBackup -Name * -FromExcServer *
```

2. Issue the database restore cmdlet.

For example:

```
Restore-DpExcBackup -Name ExchDb01 -Full  
-BACKUPDESTINATION TSM -FROMEXCSErVer PALADIN20  
-INTODB Zwen
```

3. Issue the restore cmdlets with parameter **intodb** to restore to an alternative location.

For example:

```
Restore-DpExcBackup -Name ExchDb01 -Full  
-BACKUPDESTINATION TSM -FROMEXCSErVer PALADIN20  
-OBJect 20140923100738 -INTODB ExchDb01_altRdb
```

IBM Storage Protect file space information

You might never need to know the file names or locations for your virtual machine files. However, if the underlying file structure interests you, Data Protection for VMware backups are stored under the node name of the vSphere datacenter (for example, *datacenter10*).

The following example shows the file space information for the virtual machine that is called *vm_exc10*.

```
Protect: ORION>q file datacenter10 f=d  
  
Node Name:          DATACENTER10  
Filespace Name:     \VMFULL-vm_exc10  
Hexadecimal Filespace Name:  
FSID: 61  
Collocation Group Name:  
Platform: TDP VMware  
Filespace Type: API:TSMVM  
Is Filespace Unicode?: No  
Capacity: 0 KB  
Pct Util: 0.0  
Last Backup Start Date/Time: 03/13/2014 21:29:17  
Days Since Last Backup Started: 31  
Last Full NAS Image Backup Completion Date/Time:  
Days Since Last Full NAS Image Backup Completed:  
Last Backup Date/Time From Client (UTC):  
Last Archive Date/Time From Client (UTC):  
Last Replication Start Date/Time:  
Days Since Last Replication Started:  
Last Replication Completion Date/Time:  
Days Since Last Replication Completed:  
Backup Replication Rule Name: DEFAULT  
Backup Replication Rule State: Enabled  
Archive Replication Rule Name: DEFAULT  
Archive Replication Rule State: Enabled  
Space Management Replication Rule Name: DEFAULT  
Space Management Replication Rule State: Enabled  
At-risk type: Default interval  
At-risk interval:
```

Microsoft SQL Server data protection in VMware environments

For Microsoft SQL Server workloads that are running in a VMware ESXi virtual guest machine, you can take application-consistent backups of virtual machines that are hosting Microsoft SQL Server data. You can also recover backups from the virtual machine.

Before you back up data, identify your recovery point objective (RPO). The *RPO* helps you decide how often to back up data and affects the cost that is associated with data backups.

For example, you can schedule frequent virtual machine backups for necessary recovery points. The recovery point of a virtual machine backup is the time of the backup. While change-block tracking and data deduplication offer savings, virtual machine backups can be expensive if you create and delete many virtual machine snapshots.

Most traditional in-guest data protection methods provide appropriate RPOs, but these in-guest methods lose the efficiencies that are introduced by backing up data at the virtual machine level.

You can use Data Protection for VMware and Data Protection for Microsoft SQL Server to back up data at a virtual machine level in a way that preserves backup efficiencies.

To protect Microsoft SQL Server data in a VMware environment, ensure that the following products are installed and configured on your system:

- IBM Storage Protect for Virtual Environments: Data Protection for VMware 8.1.27
- Data Protection for Microsoft SQL Server 8.1.11

For permission required to back up and restore application data for Microsoft SQL Server, see [technote 1647995](#).

Application protection is supported for VMware VMs in a VMware vSphere environment only.

Configuring the software for SQL Server data protection in a VMware environment

To protect Microsoft SQL Server workloads that are running in a VMware ESXi virtual guest machine, install and configure Data Protection for VMware. Then, install and configure Data Protection for Microsoft SQL Server.

Before you begin

The following instructions are based on a configuration scenario that might not match your environment. Adjust the configuration for your environment.

The following list summarizes the scenario for quick reference:

Windows host name

SQL10

VSS requestor node name

SQL10_VSS

Data Protection for Microsoft SQL Server node name

sql10_SQL

Virtual machine name

vm_sql10

Data mover node names

datamover10 and datamover20

Datacenter node name

datacenter10

VM file space

\VMFULL-vm_sql10

About this task

The following details describe the scenario that is used.

- A single Microsoft SQL Server database on a virtual machine that is named vm_sql10 must be recovered.
- Virtual machine vm_sql10 is protected by Data Protection for VMware by using the node name datacenter10. This node name in the IBM Storage Protect server represents the vSphere datacenter). The data mover nodes are called datamover10 and datamover20.
- The virtual machine guest is configured with the virtual machine name of vm_sql10 and the Microsoft Windows host name is SQL10.
- Data Protection for Microsoft SQL Server is installed in the guest and is configured to the IBM Storage Protect server to use node name sql10_SQL.

- The IBM Storage Protect client in the virtual guest machine is configured as the VSS requestor node and is using the node name SQL10_VSS.

Procedure

1. Follow the installation and configuration instructions that are provided with each software package.

If you install Data Protection for Microsoft SQL Server before Data Protection for VMware, you cannot specify the VMware datacenter node in the Data Protection for Microsoft SQL Server configuration wizard because the field is disabled.

2. Complete the tasks in this checklist:

- Verify that Microsoft SQL Server databases and mailboxes are hosted on VMware virtual disks.
- Verify that no Microsoft SQL Server database is hosted on raw device mapped (RDM) disks in physical compatibility mode, independent disks, or on disks that are attached directly to the guest through in-guest iSCSI.
- Verify that policies are set to keep sufficient versions of Microsoft SQL Server logs and virtual machine backups.
- Verify that SQL Server databases are on a single server and are not participating in any type of clustering, for example, failover clusters, AlwaysOn Availability Groups or AlwaysOn Failover Cluster instances.

3. Outside of the virtual guest machine, in the datamover, configure Data Protection for VMware to protect Microsoft SQL Server databases

4. Within the virtual guest machine, take the following actions:

- Verify that the Data Protection for VMware recovery agent command-line interface is configured to work with the recovery agent on the guest virtual machine.
- Configure Data Protection for Microsoft SQL Server to complete SQL Server log backups and restore SQL Server databases from a virtual machine backup.

Related tasks

[“Configuring Data Protection for VMware” on page 59](#)

You must configure Data Protection for VMware to preserve Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft Exchange Server data.

[“Configuring Data Protection for Microsoft SQL Server” on page 84](#)

After you configure Data Protection for VMware and verify that you created a virtual machine backup that is suitable for recovery of a single Microsoft SQL Server database, configure Data Protection for Microsoft SQL Server in the guest virtual machine.

Configuring Data Protection for VMware

You must configure Data Protection for VMware to preserve Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft SQL Server data.

About this task

Data Protection for VMware provides application consistency when you back up virtual machines that are hosting Microsoft SQL Servers. With these backups, you can recover the virtual machine with Microsoft SQL Server in a consistent state.

To recover only selected databases from this type of backup with IBM Storage Protect, without having to recover the entire virtual machine, preserve information about the state of the Microsoft SQL Server at the time of the virtual machine snapshot and backup. This information is collected as part of the Microsoft Volume Shadow Copy Services (VSS) interaction that occurs during a virtual machine snapshot.

For Data Protection for VMware to collect the Microsoft VSS metadata for Microsoft SQL Server, you must configure Data Protection for VMware to obtain this information from the virtual machine during the backup operation.

Procedure

1. Configure Data Protection for VMware to preserve the Microsoft VSS metadata information during a virtual machine backup for systems that are hosting Microsoft SQL Server data.

a) Locate the options file for the Data Protection for VMware data mover. On Windows systems, the options file is `dsm.opt`. On Linux systems, the options file is `dsm.sys`.

b) Specify the `INCLUDE.VMTSMVSS` option for the virtual machine.

You must set this option for virtual machine backups to preserve the Microsoft VSS metadata information. Choose from the options in the following table:

Table 8. <i>INCLUDE.VMTSMVSS</i> options	
Option	Result
<code>INCLUDE.VMTSMVSS vm_display_name</code>	When you set this option, virtual machine applications receive a notification when a backup is going to occur. This notification allows the application to commit transactions and truncate transaction logs so that the application can resume from a consistent state when the backup completes. <code>vm_display_name</code> refers to the name of the virtual machine as shown in the VMware vSphere Client and vSphere Web Client.
<code>INCLUDE.VMTSMVSS vm_display_name OPTions=KEEPSqllog</code>	When you set this option, SQL server logs are not truncated when a data mover node backs up a virtual machine that runs a SQL server. By specifying this parameter, you can manually preserve the SQL Server logs and restore SQL transactions to a specific checkpoint after the virtual machine is restored. When you specify this option, the SQL log is not truncated.

c) Verify that the Virtual Machine Disks (VMDK) that host the Microsoft SQL Server database are not being excluded from the virtual machine backup operation.

Repeat the preceding steps for all data movers that protect virtual machines that are hosting Microsoft SQL Server.

2. On each data mover, for example, `datamover10`, store the guest virtual machine credentials to Data Protection for VMware by running the following command from the IBM Storage Protect backup-archive client command line:

```
dsmc set password -type=vmguest vm_guest_display_name  
guest_admin_ID guest_admin_pw
```

This command stores the guest virtual machine credentials, which are encrypted on the system that hosts the data mover. The following minimum permissions are required for `guest_admin_ID` `guest_admin_pw`:

- Backup rights: Users with the `db_backupoperator` database role are granted to run the self-contained application data backup. If the user is a member of the SQL Server `sysadmin` fixed server role, the user can back up any databases of Microsoft SQL Server instance. The user can also back up the databases for which the user is the owner and does not have backup rights to a specific database. The guest VM user must have permission to create Volume Shadow Copies and to truncate SQL Server logs.
- Restore rights: If the database exists, you can complete the restore operation if you are a member of the `dbcreator` fixed server role, or if you are the database owner. Users with a Microsoft SQL Server `sysadmin` fixed server role have permission to restore a database from any backup sets. For other users, the situation depends on whether the database exists.

What to do next

You can verify the virtual machine backup configuration and ensure that the VMDKs are included. In addition, you can view other parameters by issuing the **backup** command with the preview option from the data mover, for example:

```
dsmc backup vm vm_display_name -preview -asnode=datacenter_node
```

You can also use the IBM Storage Protect scheduler to schedule periodic backups of your virtual machines. You can also back up the virtual machine that is hosting the Microsoft SQL Server by using the data mover command line:

```
dsmc backup vm vm_display_name -asnode=datacenter_node
```

Verifying that the configuration backs up data that can be restored

Before you can restore individual Microsoft SQL Server databases from a Data Protection for VMware virtual machine backup, you must complete at least one successful virtual machine backup. For the restore operation to work, the backup must contain Microsoft SQL Server database metadata.

Procedure

1. Issue the following data mover **query** command on one of the data mover nodes:

```
dsmc query vm vmname -detail -asnode=datacenter_node
```

where:

- *vmname* specifies the name of the virtual machine
- *datacenter_node* specifies the name of the datacenter node

2. In the command output, look for the following details:

```
Application(s) protected: MS SQL 2012 (database-level recovery)
```

Ensure that Excluded is not indicated in any Virtual Machine Disk (VMDK) status fields for virtual disks that host Microsoft SQL Server database files. The Excluded status indicates that one or more of the VMDKs that are required to recover a Microsoft SQL Server database are not being protected. For example:

```
Query Virtual Machine for Full VM backup
# Backup Date      Mgmt Class  Size      Type      A/I      Virtual Machine
-----
1 02/20/2016      STANDARD   43.94GB   IFFULL    A        vm_sql10
  12:43:59

    Size of this incremental backup: n/a
    Number of incremental backups since last full: 0
    Amount of extra data: 0
    Object fragmentation: 0
    Backup is represented by: 328 objects
    Application protection type: TSM VSS
    Application(s) protected: MS SQL 2012 (database-level recovery)
    VMDK[1]Label: Hard Disk 1
    VMDK[1]Name: [ess800_dev2] vm_sql10/vm_sql10 .vmdk
    VMDK[1]Status: Protected
    ...
    VMDK[6]Label: Hard Disk 6
    VMDK[6]Name: [ess800_dev2] vm_sql10/vm_sql10_5.vmdk
    VMDK[6]Status: Protected
```

Configuring Data Protection for Microsoft SQL Server

After you configure Data Protection for VMware and verify that you created a virtual machine backup that is suitable for recovery of a single Microsoft SQL Server database, configure Data Protection for Microsoft SQL Server in the guest virtual machine.

Procedure

1. Log on to the virtual machine that hosts the Microsoft SQL Server database.
2. Verify that the following packages are installed:
 - IBM Storage Protect recovery agent, recovery agent command-line interface (CLI), and license (from the Data Protection for VMware product package)
 - IBM Storage Protect data mover
 - Data Protection for Microsoft SQL Server

You can install the recovery agent, CLI, license, and data mover together by using the Data Protection for VMware installation program. To install the packages together, select the following advanced installation option: **Install a complete data mover for in-guest application protection**. Data Protection for Microsoft SQL Server is installed separately.

3. Configure Data Protection for Microsoft SQL Server by using the IBM Storage Protect configuration wizard. When you open the **IBM Storage Protect Node Names** page of the wizard, enter the VMware datacenter node name, Microsoft SQL Server node name, and VSS requestor node name. If the datacenter node name field is disabled, the recovery agent is not installed correctly.
4. After Data Protection for Microsoft SQL Server is configured, verify that the **Configuring Recovery Agent** rule status indicates Passed.
5. Log on to a data mover instance and complete the following steps.

Do not repeat these steps for all data mover instances.

- a) Copy the contents of the data mover options file `dsm.opt` to a temporary file named `dsm.setaccess.opt` and make the following changes to the file. Do not make these changes in the `dsm.opt` file.
 - i) Delete any line that contains an **ASNODE** entry.
 - ii) Set the **NODENAME** option to the VMware datacenter node name. For example:

```
NODENAME datacenter10
```

Tip: If the `dsm.opt` file does not contain **ASNODE** entries and the **NODENAME** option is set to the correct datacenter node, you can use this file rather than creating the `dsm.setaccess.opt` file.

- b) From the datacenter node that was defined in the **NODENAME** option, issue the **set access** command to grant the VSS requestor node access to the virtual machine backups as shown in the following example.

You must complete this step because the VSS requestor node accesses the virtual machine backups on behalf of Data Protection for Microsoft SQL Server.

If the password for the datacenter node is unknown when you run the **set access** command, you will receive an error message and the IBM Storage Protect server administrator must reset the password to issue the command.

Example

The following example shows the required parameters for the **set access** command. The parameters specify the virtual machine name (`vm_sql10`), the VSS requestor node name (`SQL10_VSS`), and the name of the options file that defines the datacenter node name (`dsm.setaccess.opt`).

```
dsmc set access backup -type=VM vm_sql10 SQL10_VSS -optfile=dsm.setaccess.opt
```

```
ANS1148I "Set Access" command successfully completed.
```

The following example shows the results of the **query access** command, which shows the backup access authorization for the VSS requestor node.

```
dsmc query access
Node name: datacenter10
Type      Node      User      Path
-----
Backup    SQL10_VSS    *        \VMFULL-vm_sql10\*\*
```

```
ANS1148I "Query Access" command completed successfully
```

Managing backups

After you configure Data Protection for Microsoft SQL Server, you can schedule backups. You must set up a virtual machine backup schedule and a Microsoft SQL Server log backup before you can start a backup.

Scheduling virtual machine backups

To ensure that your data is protected, schedule virtual machine backups.

Before you begin

By default, the maximum size allowed for a VMDK in a backup operation is 2 TB. However, the maximum is 8 TB. To increase the maximum size, use the `vmmaxvirtualdisks` option. For more information, see [Vmmaxvirtualdisks](#).

Procedure

1. Log on to the Data Protection for VMware user interface.
2. Click the **Backup** tab.
3. Click **Create Schedule** to specify a backup schedule name, source (the virtual machines to include in the backup schedule), and other scheduling options.
4. Verify that the source of the schedule includes the virtual machines that are hosting Microsoft SQL Server.
5. Verify that one of the following services is running:
 - If you are using a scheduler that is managed by a Client Acceptor Daemon (CAD), ensure that the CAD service is running on the data mover.
 - If you are using the stand-alone scheduler, ensure that the scheduler service is running.

Scheduling Microsoft SQL Server log backups

After the virtual machine backup schedule is created, you can create the Microsoft SQL Server log backup schedule.

About this task

Backing up SQL server logs provides a more granular level of recovery points. You might find it unnecessary to back up SQL server logs if the frequency of your backups provides you with enough recovery points, and assuming that you did not specify the `INCLUDE.VMTSMVSS vm_display_name` `OPTIONS=KEEPSqllog` option for the backup.

Procedure

1. Start the Data Protection for Microsoft SQL Server user interface from the virtual machine that is hosting Microsoft SQL Server.
2. In the navigation pane, expand the **Manage** node.
3. Under the **Manage** node, right-click **Scheduling** > **Scheduling Wizard**.
4. Open the **Scheduling Wizard** to identify the schedule name and time.
5. For the **Define the Scheduled Task** page, select **Command Line**.
6. Click the icon to select the SQL Server template. Click **Next**.
7. Use the command-line interface and SQL Server template to specify the database log backup, for example:

```
tdpsqlc backup * log /truncate=yes 2>&1
```

Tip: Alternatively, you can schedule Microsoft SQL Server backups by using the IBM Storage Protect centralized scheduling service. This service helps you to create a backup schedule for all Microsoft SQL Server instances on a virtual machine.

Verifying backups

After you create a backup, verify that you can query the virtual machine backups and the database backups from the Data Protection for Microsoft SQL Server interface.

About this task

You can recover one or more Microsoft SQL databases based on your recovery point objectives.

Procedure

1. From Microsoft Management (MMC), select a Microsoft SQL Server.
2. Click the **Recover** tab.
3. Select **View** > **Databases**. A list of Microsoft SQL Server database backups that can be restored is displayed.

Microsoft SQL Server databases that are backed up with Data Protection for VMware are identified with the backup method *vmvss*. Microsoft SQL Server logs that are backed with Data Protection for Microsoft SQL Server are identified with the backup method *Legacy*.

Managing versions of backups

By using Data Protection for Microsoft SQL Server, you can manage expiration of backups. You can specify the number of snapshot backups to retain and the length of time to retain snapshots.

About this task

To set the retention for Microsoft SQL Server backups, complete the following steps. This procedure assumes that you want to retain backups for 30 days.

Procedure

1. Define the retention parameters in the management class that is used for virtual machine backups. For example:

```
Retain extra versions = 30  
Retain only versions = 30  
Versions data exists = nolimit  
Versions data deleted = nolimit
```


Use the `vmmc` option in the data mover option file to specify the management class that is used for the virtual machine backups.

Scheduled virtual machine backups are associated with the Data Protection for VMware client.

2. Define the retention parameters in the management class that is used for Microsoft SQL Server backups.

For example:

```
Retain extra versions = 0
Retain only versions = 1
Versions data exists = nolimit
Versions data deleted = nolimit
```

Specify the management class for the Microsoft SQL Server backups in the `dsm.opt` file that is used by the Data Protection for Microsoft SQL Server agent. See the following `INCLUDE` options:

```
INCLUDE *:...\*log management_class_name
INCLUDE *:...\log\..\* management_class_name
```

3. With Data Protection for Microsoft SQL Server running on the virtual machine, issue the **inactivate** command to explicitly deactivate all active log backups for all databases on the Microsoft SQL Server. For example:

```
tdpsqlc inactivate * log=* /OLDERTHAN=30
```

Log backups that are created by Data Protection for Microsoft SQL Server must be explicitly deactivated because the full database backups are being completed by Data Protection for VMware. This configuration allows for a one-day grace period after the Microsoft SQL Server log backups are deactivated before the IBM Storage Protect server deletes them.

Tip: You can retain log backups on the server only if the full database backup with which they are associated are retained. In the management class, set the **RETONLY** value for log backups to match the **RETEXTRA** parameter for full database backups.

Verifying that virtual machine backups do not exclude Microsoft SQL Server volumes

The volumes in Virtual machine disks (VMDKs) must contain the Microsoft SQL Server databases that are not excluded from the Data Protection for VMware backup processing.

About this task

The databases cannot be on physical compatibility mode raw device mapping (RDM) disks, independent disks, or on disks that are attached directly to the guest operating system through iSCSI.

Procedure

1. Ensure that any `EXCLUDE.VMDISK` statements in the Data Protection for VMware data mover that is used to back up the virtual machine do not inadvertently exclude VMDKs that are hosting volumes that contain Microsoft Exchange Server files, file space, database, and mailboxes.

For example:

- `vm_sql10.vmdk` contains logical volume C:
 - `vm_sql10.vmdk` contains logical volumes E: and F:
 - The label for `vm_sql10_1.vmdk` is *Hard Disk 1*.
 - The label for `vm_sql10_2.vmdk` is *Hard Disk 2*.
 - The Microsoft SQL Server database files to be backed up are on the E: and F: drive.
2. Verify that no statements exclude `vm_exc10_2.vmdk` from the virtual machine backup by ensuring that the data mover does not contain the following or similar statements:

```
EXCLUDE.VMDISK VM_SQL10 "Hard Disk 2"  
EXCLUDE.VMDISK * "Hard Disk 2"
```

Alternatively, if you exclude most hard disks, you must explicitly include the virtual machine disks by using one of the following statements:

```
INCLUDE.VMDISK VM_SQL10 "Hard Disk 2"  
INCLUDE.VMDISK * "Hard Disk 2"
```

Include and exclude statements are processed from bottom to top as they are displayed in the `dsm.opt` file. To achieve the goal, enter the statements in the correct order.

You can specify the exclusion and inclusion of a virtual machine disk from the command-line interface:

```
dsmc backup vm "VM_SQL10:-vmdisk=Hard Disk 2" -asnode=datacenter10
```

Restoring data

After you back up data, you can recover the data based on a recovery point objective (RPO).

A recovery operation restores a full backup of the Microsoft SQL Server database from the Data Protection for VMware backup.

If you restore the entire virtual machine, all Microsoft SQL Server databases on the virtual machine are restored and recovered to the point of the virtual machine backup. In this scenario, you cannot restore and recover any backups that were created after that point.

Starting the Microsoft iSCSI Initiator Service

The iSCSI protocol is used to mount the disks that are used for a recovery operation. Ensure that the Microsoft iSCSI Initiator Service is started and is set to the automatic startup type on the system where the data is to be restored.

Procedure

1. In the Windows **Services** list, right-click **Microsoft iSCSI Initiator Service**.
2. Click **Properties**.
3. On the **General** tab, set the following options:
 - a) In the **Startup type** list, select **Automatic**.
 - b) Click **Start**, and then click **OK**.

Results

In the **Services** list, **Microsoft iSCSI Initiator Service** shows a status of **Started** and the startup type is **Automatic**.

Restoring database backups by using the graphical user interface

You can recover a full Microsoft SQL Server database backup from a virtual machine backup by using the Data Protection for Microsoft SQL Server graphical user interface.

Procedure

1. To start a full database recovery from a virtual machine, start Microsoft Management (MMC). In the navigation pane, expand the **Protect and Recover** node and select a Microsoft SQL Server server.
2. On the **Recover** tab, select **Database Restore**. All backups, including all database backups from a virtual machine backup, are listed.
3. Select a full database backup to restore.
4. In the **Actions** pane, click **Restore**.

Restoring data by using the command-line interface

If you prefer, you can use the command-line interface to start a full Microsoft SQL Server database recovery from a virtual machine.

Procedure

1. Issue the **query** command to find the full and log database backups.

The following example finds all backups for the Microsoft SQL Server database called sql_db10.

```
tdpsqlc q tsm sql_db10
IBM Storage Protect for Databases:
Data Protection for Microsoft SQL Server
Version 8, Release 1, Level 0.0
...
Querying IBM Storage Protect Server for Backups ....
Backup Object Information -----
SQL Server Name ..... SQL10
SQL Database Name ..... sql_db10
Backup Method ..... VMVSS
...
Backup Creation Date / Time ..... 11/14/2014 13:41:18
...
Backup Object Information
-----
SQL Server Name .....
SQL10 SQL Database Name .....sql_db10
Backup Method ..... LgcY
...
Backup on Secondary Replica .....
No Backup Object State .....
Active Backup Creation Date / Time ..... 11/14/2014 15:46:07
...
The operation completed successfully. (rc = 0)
```

2. To restore the database without applying transaction logs, issue the database **restore** command as shown in the following example:

```
tdpsqlc restore databaseName /backupMethod=vmvss
```

The following examples show the output of the command when you specify the Microsoft SQL Server database called sql_db10.

```
tdpsqlc restore sql_db10 /backupmethod=vmvss /sqlserver=sql10
/fromsqlserver=sql10 /recovery=no
IBM Storage Protect for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2016. All rights reserved.

Connecting to SQL Server, please wait...
Querying IBM Storage Protect Server for Backups ....
Connecting to IBM Storage Protect Server as node "SQL10_SQL"...
Connecting to Local DSM Agent "SQL10"...
Using backup node "SQL10_SQL"...
Starting Sql database restore...

Beginning VSS restore of "sql_db10"...

Restoring "sql_db10" via file-level copy from snapshot(s). This
process may take some time. Please wait

Files Examined/Completed/Failed: [ 2 / 2 / 0 ] Total Bytes: 3146070

VSS Restore operation completed with rc = 0
Files Examined : 2
Files Completed : 2
Files Failed : 0
Total Bytes : 3146070
Total LanFree Bytes : 0

The operation completed successfully. (rc = 0)
```

3. After the full database restore operation is completed successfully, issue the command to restore the logs.

For example, to restore all logs based on the restored Microsoft SQL database `sql_db10`, issue the following command.

```
tdpsqlc restore databaseName /backupMethod=vmvss  
/recovery=no
```

You can also use the `/stopat` option to specify a more granular point in time.

```
tdpsqlc restore sql_db10 log=* /sqlserver=sql10  
/fromsqlserver=sql10 /recovery=yes  
IBM Storage Protect for Databases:  
Data Protection for Microsoft SQL Server  
Version 8, Release 1, Level 0.0  
(C) Copyright IBM Corporation 1997, 2016. All rights reserved.  
  
Connecting to SQL Server, please wait...  
Starting Sql database restore...  
Connecting to IBM Storage Protect Server as node "SQL10_SQL"...  
Querying IBM Storage Protect server for a list  
of database backups, please wait...  
  
Beginning log restore of backup object sql_db10\20131114154607\00000DB0,  
1 of 3, to database sql_db10 ...  
  
Beginning log restore of backup object sql_db10\20131114155130\00000DB0,  
2 of 3, to database sql_db10 ....  
  
Total database backups inspected: 3  
Total database backups requested for restore: 3  
Total database backups restored: 3  
Total database skipped: 0  
Throughput rate: 134.32 Kb/Sec  
Total bytes transferred: 385,536  
Total LanFree bytes transferred: 0  
Elapsed processing time: 2.80 Secs  
The operation completed successfully. (rc = 0)
```

What to do next

You can restore inactive backups by using the Data Protection for Microsoft SQL Server command-line interface, **TDPSQLC**. When you issue the **restore** command, specify the database object name for the specific backup.

To obtain the database object name, issue the following command:

```
tdpsqlc q tsm dbname full /all
```

After you have the database object name value, specify the database object name on the `/Object=objectname` parameter of the **TDPSQLC restore** command, where *objectname* is the database object name. For example:

```
tdpsqlc restore db44 /object=20140311131051 /backupdestination=tsm  
/backupmethod=vmvss
```

Restriction: You cannot recover a Microsoft SQL database to an alternative location on the virtual machine.

Restoring Microsoft SQL Server log backups

After the full database is restored successfully, you can restore transaction logs.

Procedure

1. Select a Microsoft SQL Server, and click the **Recover** tab.
2. Verify that the **AutoSelect** option is set to **False**.
3. Change the **RunRecovery** option to **True**.

4. Select all the logs that you want to recover.
5. Click **Restore**.

Restoring relocated and deleted mailboxes

The backup solution for restoring databases and log files that are relocated and deleted after a virtual machine backup consists of Data Protection for VMware and Data Protection for Microsoft SQL Server.

Before you begin

Decide where the database and log file data is to be restored.

About this task

When you restore the backups, and complete a full database restore operation from the backup, Data Protection for VMware restores the files to their original location.

If database or log files are relocated during the backup cycle, Data Protection for Microsoft SQL Server restores the files in their original locations.

If any databases or log files were created during the backup cycle, Data Protection for Microsoft SQL Server re-creates the new files. If database or log files were deleted during the backup cycle, those files are not restored.

Procedure

1. Use Data Protection for VMware to back up the virtual machine.
Consider the following example. You back up virtual machine `vm_sql110` that includes Microsoft SQL Server database `moose` at 2:00 p.m. The Microsoft SQL Server database consists of the following files at 2:00 p.m:
 - `C:\sql dbs\moose\moose.mdf`
 - `C:\sql dbs\moose\moose_log.ldf`
2. Relocate a database backup to an alternate location.
Consider the following example. You want to relocate the database `moose` at 6:00 p.m. to the following location:
 - `E:\sql dbs\moose\moose.mdf`
 - `F:\sql dbs\moose\moose_log.ldf`
3. Add files to the database backup.
Consider the following example. You want to add two new files to database `moose` at 7:00 p.m. The database now consists of the following files:
 - `E:\sql dbs\moose\moose.mdf`
 - `F:\sql dbs\moose\moose_log.ldf`
 - `E:\sql dbs\moose\moose2.ndf`
 - `F:\sql dbs\moose\moose2_log.ldf`
4. Use Data Protection for Microsoft SQL Server to complete a log backup.
Consider the following example. You start a log backup at 9:00 p.m.
5. Restore the database backup.
Consider the following example. You want to restore the entire `moose` database.
 - You restore the full database from the Data Protection for VMware backup with **`runrecovery=false`**.
 - At 9:00 p.m, you restore the log backup and apply it.The `moose` database is restored to the following location:
 - `C:\sql dbs\moose\moose.mdf`

- C:\sql dbs\moose\moose_log.ldf
- E:\ sql dbs\moose\moose2.ndf
- F:\ sql dbs\moose\moose2_log.ldf

The full virtual machine restore restores the files to their original location. When you applied the log backup, the files that were added after the relocation are restored.

Sample script for validating full virtual machine backups

Before you back up Microsoft SQL Server logs, verify that you have a valid full virtual machine backup. One procedure for checking for the existence of a full virtual machine backup is to schedule the usage of a script.

This sample script checks for the instance of a full backup and then runs the Microsoft SQL Server log backups if a full virtual machine backup exists. This script can be used with a scheduler service such as the IBM Storage Protect scheduler.

```
@echo off
dsmc q vm sql01_SQL -detail -asnode=datacenter01 | find /c
"database-level recovery" > c:\temp.txt
SET /p VAR=<c:\temp.txt

if %VAR% == "1" (
tdpsqlc back * log
) ELSE (
echo "There is no full backup"
set ERRORLEVEL=1
)
```

This script produces the following output:

```
IBM Storage Protect for Databases:
Data Protection for Microsoft SQL Server
Version 8, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2016. All rights reserved.
Connecting to SQL Server, please wait...
Starting SQL database backup...
Connecting to IBM Storage Protect Server as node 'SQL01_SQL'...
Using backup node 'SQL01_SQL...'
AC05458W The IBM Storage Protect Server 'backup delete' setting for node (SQL01_SQL)
is set to NO. It should be set to YES for proper operation. Processing will continue.
Beginning log backup for database model, 1 of 2.
Full: 0 Read: 87808 Written: 87808 Rate: 32.54 Kb/Sec
Database Object Name: 20140303011509\000007CC
Backup of model completed successfully.
Beginning log backup for database sql db test2, 2 of 2.
Full: 0 Read: 88832 Written: 88832 Rate: 132.44 Kb/Sec
Database Object Name: 20140303011511\000007CC
Backup of sql db test2 completed successfully.
Total SQL backups selected: 4
Total SQL backups attempted: 2
Total SQL backups completed: 2
Total SQL backups excluded: 2
Total SQL backups deduplicated: 0
Throughput rate: 51.85 Kb/Sec
Total bytes inspected: 176,640
Total bytes transferred: 176,640
Total LanFree bytes transferred: 0
Total bytes before deduplication: 0
Total bytes after deduplication: 0
Data compressed by: 0%
Deduplication reduction: 0.00%
Total data reduction ratio: 0.00%
Elapsed processing time: 3.33 Secs
The operation completed successfully. (rc = 0)
```

You can also use the IBM Storage Protect activity log and extended summary table to determine whether virtual machine backups are successful.

IBM Storage Protect file space information

You might never need to know the file names or locations for your virtual machine files. However, if the underlying file structure interests you, Data Protection for VMware backups are stored under the node name of the vSphere datacenter (for example, datacenter10).

The following example shows the file space information for the virtual machine that is called vm_sql110.

```
Protect: ORION>q file datacenter10 f=d

Node Name:          DATACENTER10
Filespace Name:      \VMFULL-vm_sql110
Hexadecimal Filespace Name:
FSID: 61
Collocation Group Name:
Platform: TDP VMware
Filespace Type: API:TSMVM
Is Filespace Unicode?: No
Capacity: 0 KB
Pct Util: 0.0
Last Backup Start Date/Time: 03/13/2014 21:29:17
Days Since Last Backup Started: 31
Last Full NAS Image Backup Completion Date/Time:
Days Since Last Full NAS Image Backup Completed:
Last Backup Date/Time From Client (UTC):
Last Archive Date/Time From Client (UTC):
Last Replication Start Date/Time:
Days Since Last Replication Started:
Last Replication Completion Date/Time:
Days Since Last Replication Completed:
Backup Replication Rule Name: DEFAULT
Backup Replication Rule State: Enabled
Archive Replication Rule Name: DEFAULT
Archive Replication Rule State: Enabled
Space Management Replication Rule Name: DEFAULT
Space Management Replication Rule State: Enabled
At-risk type: Default interval
At-risk interval:
```

Application protection for Active Directory domain controllers

Data Protection for VMware provides back up and restore protection for VMs that host Microsoft Active Directory Domain Controllers in both stand-alone and clustered environments. A clustered environment contains multiple domain controllers that participate in Active Directory.

Non-authoritative restore recovers the Active Directory (or domain controller) to the version taken at the time of the backup. When the recovered Active Directory (or domain controller) is restored, it is updated with information from the other domain controllers through the existing replication process.

Environment requirements

Data Protection for VMware protects Windows VM guests that host Active Directory Domain Controllers. The following guest versions that host Active Directory Domain Controllers are supported:

- **Windows** Microsoft Windows Server 2012
- **Windows** A current version of VMware Tools must be installed and must be running on the VM guest at the time that it is backed up. This VM guest must be powered on for Data Protection for VMware to detect Active Directory. Otherwise, Active Directory will not be detected and restore protection will be unavailable.
- **Windows** Active Directory on Microsoft Windows Server versions 2016 and later require no special support.

Restriction:

When a VM guest contains Active Directory or a domain controller, ensure that Windows NT Directory Services (NTDS) is running so that the VSS backups and domain controller discovery can function correctly. You cannot use application protection for domain controllers to complete these tasks:

- Run a file restore of Active Directory objects
- Back up and restore VMs that run Active Directory Lightweight Directory Services (AD LDS)
- Recover expired Active Directory tombstone objects

Tip: To help prevent Active Directory objects from expiring, run backups more frequently than the default tombstone life of 60 days.

- Run a full VM instant restore operation

Chapter 5. Data Protection for VMware commands and options

Data Protection for VMware provides command-line interfaces (CLIs) that you can use as alternatives to the graphical user interfaces (GUIs) and option files that are provided with the product.

The primary CLI for Data Protection for VMware is run from the **dsmc** command. This CLI provides commands and options that you can use to manage virtual machines (VMs) that are in a vSphere environment.

A secondary CLI is available for troubleshooting problems with the Data Protection for VMware vSphere GUI. This CLI is run from the **vmcli** command.

A CLI is also available for the IBM Storage Protect recovery agent.

Restriction: When running commands with parameters that contain an ampersand (&), it must be enclosed in either single or double quotation marks. For example, **-optfile="dsm.VM14_ '&'1.opt"**

dsmc command-line interface

This is the primary CLI for use with Data Protection for VMware. This CLI must be run on a system that contains the data mover.

Related information

[Using commands](#)

dsmc commands

The following dsmc commands are available to back up, restore, and configure VMs in your vSphere environment.

For information about the commands, click the following links:

- [Backup VM](#)
- [Delete Backup](#)
- [Expire](#)
- [Query VM](#)
- [Restore VM](#)
- [Set Access](#)
- [Set Password](#)
- [Set Vmtags](#)

dsmc command options

The following options are available for use with specific dsmc commands. You can use these options to refine the operation of a command.

For information about the options, click the following links:

- [Csv](#)
- [Domain.vmfull](#)
- [Virtual machine exclude options](#)
- [Exclude.vmdisk](#)
- [Virtual machine include options](#)
- [Include.vm](#)
- [Include.vmdisk](#)
- [INCLUDE.VMSNAPSHOTATTEMPTS](#)

[INCLUDE.VMTSMVSS](#)
[Mbobjrefreshthresh](#)
[Mbpctrefreshthresh](#)
[Mode](#)
[Vmautostartvm](#)
[Vmbackdir](#)
[Vmbackuplocation](#)
[Vmbackupmailboxhistory](#)
[Vmbackuptype](#)
[Vmchost](#)
[Vmcpw](#)
[Vmctlmc](#)
[Vmcuser](#)
[Vmdatastorethreshold](#)
[Vmdefaultdvportgroup](#)
[Vmdefaultdvswitch](#)
[Vmdefaultnetwork](#)
[Vmdiskprovision](#)
[Vmenabletemplatebackups](#)
[Vmexpireprotect](#)
[Vmiscsiadapter](#)
[Vmiscsiserveraddress](#)
[Vmlimitperdatastore](#)
[Vmlimitperhost](#)
[Vmmaxbackupsessions](#)
[Vmmaxparallel](#)
[Vmmaxrestoresessions](#)
[Vmmaxvirtualdisks](#)
[Vmmc](#)
[Vmmontage](#)
[Vmnoprmdisks](#)
[Vmnovrmdisks](#)
[Vmpreferdagpassive](#)
[Vmprocessvmwithindependent](#)
[Vmprocessvmwithprdm](#)
[Vmrestoretype](#)
[Vmskipctlcompression](#)
[Vmskipmaxvirtualdisks](#)
[Vmstoragetype](#)
[Vmtagdefaultdatamover](#)
[Vmtagdatamover](#)
[Vmtempdatastore](#)
[Vmverifyifaction](#)
[Vmverifyiflatest](#)
[Vmvstorcom](#)
[Vmvstortransport](#)
[Vmtimeout](#)

vmcli command-line interface

This is a secondary CLI that provides commands and options that you can use to troubleshoot problems with the Data Protection for VMware vSphere GUI.

About this task

Windows | **Linux** The following commands are available:

[“Backup” on page 97](#)

Initiate® full and incremental backups of your VMs.

[“Restore” on page 99](#)

Restore backups of your VMs.

[“Inquire_config” on page 104](#)

View configuration information about the backup database.

[“Inquire_detail” on page 106](#)

View configuration information about the backup environment.

[“Set_domain” on page 108](#)

Apply changes to the domain settings.

[“Set_option” on page 109](#)

Set a parameter in the vmcliprofile.

[“Set_password” on page 110](#)

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

[“Get_password_info” on page 112](#)

View the status of guest credentials that are set for the managed data centers.

[“Start_guest_scan” on page 113](#)

Scan guest VMs for application information.

Example

Access the CLI in the following directories:

Linux
/opt/tivoli/tsm/tdpvmware/common/scripts

Windows (64-bit)
C:\Program Files\IBM\StorageProtect\Framework\VEGUI\scripts

For CLI messages that contain the FMM prefix, message information is available in the IBM Documentation:

[FMM, FMF, FMV, FMX, FMY: IBM Storage Protect Snapshot messages](#)

Backup

Use this **vmcli** command to start IFFULL and IFINCREMENTAL backups of your VMs or VM templates.

Syntax

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

```
vmcli -f backup -t backupType -I backupObjectListFile -d datacenternodename  
| providervDCnodename -o datamovernodename [--name taskName] [--description  
descriptionInFile.txt] [--s tsmserverhostname] [--n vctrclinodename] [--p tsmserverport] [--w  
vmBackupLocation]
```

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

Parameters

Before you issue 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.

When a backup operation is running, there is no command or method available to stop the backup, including the Ctrl + C command. You must wait for the operation to complete on its own.

The **vmcli -f backup** command requires that the **VE_VCENTER_NODE_NAME** is set correctly in the **vmcli** profile. You cannot overwrite this parameter with a command-line entry.

The data mover system (the vStorage Backup Server where the IBM Storage Protect backup-archive client is installed) must not set the **ASNODENAME** option.

-t backupType

Specify the type of backup to complete. You can choose from one of the following types:

TSM_IFFULL

Creates an incremental forever full backup of the specified backup objects. When **TSM_IFFULL** is specified, template VMs that are unchanged since the last backup are also included.

TSM_IFINCR

Creates an incremental forever incremental backup of the specified backup object. This type backs up only the changed data since the last backup. This type is the default.

The backup process does not create a snapshot of template VMs in the same manner that a snapshot is created for regular VMs. As a result, VMware VDDK advanced transports (SAN, HotAdd, and NBDSSL modes), change block tracking (CBT), and incremental backups are not available.

-I backupObjectListFile

Specify the file that contains the list of objects to back up. Each line contains one specification for backup.

In vSphere mode, 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
```

Restrictions:

- Do not specify a VM host name in the *backupObjectListFile*. Data Protection for VMware does not support backing up a VM that is identified by the VM host name.
- When you specify the name of a VM by 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 that is used in the VM name. Therefore, use caution when you specify keyword values. In addition, backing up a VM that contains a comma in its name is not supported.
- Data Protection for VMware support for VM backup operations is limited to VM names and datacenter names that contain English 7-bit ASCII characters only. VM names and datacenter names that use other language characters are not currently supported. More character restrictions are listed in [Appendix A, “Troubleshooting,”](#) on page 163.
- A VMware vCenter allows the existence of two VMs with the same name. However, Data Protection for VMware does not support backing up two VMs with the same name. To prevent errors or backup failures, do not have two VMs with the same name in a vCenter.

-d datacenternodename | providervDCnodename | organizationvDCnodename

When the **VE_TSM_MODE** parameter specifies **VSPHERE**, specify the datacenter node name.

-o *datamovernodename*

Specify the data mover node name. This name is the node name for the data mover 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 *tsmserverhostname*]

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

[-n *vmclinodename*]

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

[-p *tsmserverport*]

Specify the port of the IBM Storage Protect 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.

[-w *vmBackupLocation*]

Specifies the location for a VM backup. The location determines whether the VM is backed up on the IBM Storage Protect server; is persisted as a snapshot on the hardware storage; or both. You can specify one of the following values:

SERVER

VMs are backed up to the IBM Storage Protect server. This value is the default.

LOCAL

VMs are backed up on the hardware storage. The backup is a full VM image snapshot, even if the **-t *backupType*** parameter specifies an incremental backup.

To create a local backup, the VM must be stored in a VMware virtual volume (VVOL) datastore. If any virtual disk of the VM is not in a VVOL datastore, the local backup is not allowed.

Because no network data movement is needed for local snapshots, backup and restore operations can be faster than server backup and restore operations. By restoring from a local snapshot, you can only revert an existing VM to an earlier point in time. You also cannot restore a deleted VM, and you cannot restore a VM to a different name or location.

BOTH

VMs are backed up to the IBM Storage Protect server and are also backed up locally. If the **-t *backupType*** parameter specifies an incremental backup, this setting applies only to the server backup. The local backup is always a full VM image backup.

Restore

Use this **vmcli** command to restore backups of your VMs or VM templates.

Syntax

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

```
vmcli -f restore -I restoreObjectListFile -d datacenternodename -o datamovernodename [-s tsmserverhostname] [-n vmclinodename] [-p tsmserverport] [-vmrestoretype (noninstant | instantrestore | instantaccess | mount | vmcleanup | vmfullcleanup | mountcleanup)] [-w vmBackupLocation]
```

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

Parameters

The **vmcli -f restore** command requires that the **VE_VCENTER_NODE_NAME** is set correctly in the `vmcliprofile`. You cannot overwrite this parameter with a command-line entry.

The data mover system (the vStorage Backup Server where the data mover is installed) must not set the `ASNODENAME` option.

-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 keyword:

backupid

Each line must begin with the backupid. The syntax is `backupid:your_backup_ID`. Specify the IBM Storage Protect Object ID for a specific VM backup. Locate the Object ID by using the **vmcli -f inquire_detail** command. This keyword is required for a restore operation.

The *restoreObjectListFile* uses the following keywords:

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: Restrictions: When you specify a keyword in the *restoreObjectListFile*, Data Protection for VMware does not differentiate between a colon (:) used as a keyword separator or a colon that is used in a keyword value. Therefore, use caution when you specify keyword values. In addition, Data Protection for VMware support for VM restore operations is limited to VM names and VMware datacenter names that contain English 7-bit ASCII characters only. VM names and datacenter names that use other language characters are not currently supported. Additional character restrictions are listed in [Appendix A, “Troubleshooting,” on page 163](#).

vmname

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

-vmdk=cnfg

Specify that the virtual machine configuration information is restored. The configuration information is always restored when the entire virtual machine is restored. However, by default the configuration is not restored when you restore only selected disks with the **vmdk=disk label** option.

Ordinarily, restoring configuration information to an existing virtual machine fails because the restored configuration information conflicts with the existing virtual machine configuration information. Use this option if the existing configuration file for a virtual machine on the ESX server was deleted, and you want to use the backed up configuration to re-create it.

For example, this entry in *restoreObjectListFile* restores all VMDKs for virtual machine VM1 and keeps the same name:

```
backupid:26801107 vmname:VM1:-vmdk=cnfg
```

vmdk=disk label

If Change Block Tracking (CBT) is enabled on the target VM, SAN transport is not supported. If a VM with multiple disks has been backed up using Data Protection for VMware and a restore is done from one or more disks to the existing VM using SAN transport method, this restore will fail if Change Block Tracking (CBT) is enabled for this VM.

In vSphere mode, specify the disk label of a virtual disk to include in the restore operation. Use this parameter only if you want to restore one or more specific disks, but not all disks. Repeat this parameter for each disk that you want to restore.

For example, this entry in *restoreObjectListFile* restores the VMDKs named Hard Disk 1 and Hard Disk 2 as a new virtual machine.

```
backupid:26801107
vmname:myvm:vmdk=Hard Disk 1:vmdk=Hard Disk 2::vmname:newname
```

The following considerations apply to each disk that you want to restore:

- – The disk must exist on the VM before you initiate the restore operation. If the disk does not exist, you must create it. You can run the `dsmc dsmc restore vm` command with the **-preview** parameter to identify the original disk label, capacity, and datastore. The **-preview** output does not include provisioning information.
- The existing disk must be at least as large as the disk you want to restore.
- The existing disk label must be the same as the disk you want to restore.
- Any data on the existing disk is overwritten.

Only the specified disks are restored. Other disks on the VM are not altered.

The VM that you are restoring the disk to must be powered off before you initiate the restore operation.

-vmdk=disk label

Specify the disk label of one or more virtual disks to exclude from the restore operation.

For example, this entry in *restoreObjectListFile* restores all VMDKs except the one named Hard Disk 1 as a new virtual machine:

```
backupid:26801107
vmname:myvm:-vmdk=Hard Disk 4::vmname:newname
```

This entry restores VMDKs for the virtual machine as a new virtual machine without configuration information:

```
backupid:26801107
vmname:oldvmname:-vmdk=cnfg::vmname:newname
```

newdatacentername

When you want the restore destination to be a different datacenter, specify the name of that datacenter 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 VMware datastore where the VM is to be restored.

For example, a datastore name such as `datastore1` is supported. A datastore URL such as `sanfs://vmfs_uuid:4d90pa2d-e9ju45ab-065d-00101a7f1a1d/` is not supported. The datastore can be on a SAN, NAS, iSCSI device, or VMware virtual volume (vVol).

vmtempdatastore

When you want to issue an instant restore operation, specify a temporary datastore on the ESX host. This temporary datastore contains the configuration information and data of the VM that is created during the operation.

vmautostartvm

When a VM is created for instant access (**vmrestoretype instantaccess**), specify whether to automatically start the VM:

YES

The VM created for instant access is automatically started.

NO

The VM created for instant access is not automatically started. It must be manually started by the user. This value is the default.

vmdiskprovision

Specify the type of provisioning for the VM disk that is restored during an instant restore process (**vmrestoretype instant**):

THICK

The disk is created with thick provisioning. This value is the default.

THIN

The disk is created with thin provisioning.

An example *restoreObjectListFile*:

```
# 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 vmtempdatastore:datastore2temp  
vmdiskprovision:thin
```

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.

The *restoreObjectListFile* uses the following keywords for mount operations:

vmostype

Specify the type of operating system for the backed up VM.

AUTOMATIC

The operating system of the backed up VM is automatically detected. This value is the default.

LINUX

The operating system of the backed up VM is Linux.

WINDOWS

The operating system of the backed up VM is Windows.

exportfs

Exports the mounted file system to the location specified by the value of the **exportparameter**.

YES

The mounted file system is exported.

NO

The mounted file system is not exported. This value is the default.

exportparameter

The location where the file system is exported.

Linux **IP or machine name**

The IP address or name of the machine that mounts the exported file system.

Windows **user name**

The user name that is allowed to access Windows Share. It is the user's responsibility to be aware of which users and groups have access to their shared files.

mountpoint mount point path

Specify the path of the mount point.

Linux The default value is `/mnt/vmname`.

Windows The default value is `D:\tsmvemount\vmname`.

mounttag string

This string is text that you enter to make the mount point name easier to identify when you search on the local file system. Specify this string as part of the mount path.

Linux The full path to a disk is `/mount root/tag/vmname/snapshot date and time/file system number`. For example:

```
/mnt/ticket1-4711/VM1/2013-12-12-12:12:12/disk1
```


Windows The full path to a disk is `mount root\tag\vmname\snapshot date and time\file system number`. For example:

```
C:\Users\Admin\ticket-4711\VM1\2013-12-12-12:12:12\disk1
```

An example `restoreObjectListFile` for mount operations:

Linux

```
backupid:1167852 vmname:VM-Lin4 mounttag:limor exportparameters:9.123.456.78
exportfs:yes vmostype:linux mountpoint:/tmp/tsm-mounts
```

Windows

```
backupid:1167850 vmname:VM-Name3 mounttag:limor exportparameters:WinUser1
exportfs:yes vmostype:windows mountpoint:C:\temp\mnt
```

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 vmtempdatastore:datastore2temp
vmdiskprovision:thin
```

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 106](#) provides current configuration information about the backup environment.

-d datacenternodename

Specify the datacenter node name.

-o datamovernodename

Specify the data mover node name. This name is for the backup-archive client node 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 IBM Storage Protect server. If this parameter is not specified, the value in the profile is used.

[-n vmclinodename]

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

[-p tsmserverport]

Specify the port of the IBM Storage Protect 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.

Windows **[-vmrestoretype (noninstant | instantrestore | instantaccess | mount | vmcleanup | vmfullcleanup | mountcleanup)]**

In a vSphere environment, specify this option to switch between the following operations: existing restore, instant access, or instant restore. Instant access and instant restore capability is supported only for VMware VMs that are hosted on VMware ESXi 5.1 servers, or later versions. The **vmrestoretype** parameter uses the following keywords:

noninstant

A full VM restore is issued.

instantrestore

The VM is started during the restore process.

instantaccess

The VM might be started but it is not restored.

mount

The volumes of the VM defined in the input file are mounted in read-only mode on the data mover. On Linux, all the volumes of the VM are mounted as a Network File System (NFS). On Windows, all the volumes of the VM are mounted as a Common Internet File System (CIFS).

vmcleanup

Components that are no longer needed are cleaned up.

vmfullcleanup

The VM and all its components are cleaned up, regardless of the current state.

mountcleanup

All mounted volumes of the selected VM are cleaned up. This cleanup task includes removing file systems that were exposed for the restore operation and the file shares (CIFS, NFS).

Restriction: When an instant restore or instant access operation that is issued from the data mover (**dsmc**) is followed by an instant restore or instant access operation that is issued from the Data Protection for VMware command-line interface (**vmcli**) or Data Protection for VMware vSphere GUI, the TDPVMwareMount service must be restarted. This situation applies only when the Data Protection for VMware command-line interface accesses the IBM Storage Protect server with a node name different from the one used by the data mover. This restriction applies to any order of operations between the two products.

Restart the service by going to **Start > Administrative Tools > Computer Management > Services and Applications > Services**. Look for service name IBM Storage Protect recovery agent in the **Services** window. The path to the **Services** window might vary depending on your operating system.

The service does not have to be restarted when the VMware datacenter name is specified with the **asnodename** option in the **dsm.opt** file.

-w vmBackupLocation

Specifies the location of the backups from which to restore the VM. You can specify one of the following values:

SERVER

The VM is restored from the IBM Storage Protect server. This value is the default.

LOCAL

The VM is restored from a persisted snapshot on the hardware storage. By restoring from a local snapshot, you can only revert an existing VM. You cannot restore a deleted VM, and you cannot restore a VM to a different name or location.

This parameter is not valid when the **vmrestoretype** parameter is also specified, unless the **vmrestoretype** is set to **mountcleanup**.

Inquire_config

Use this **vmcli** command to view configuration information about the IBM Storage Protect nodes associated with Data Protection for VMware.

Syntax

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

```
vmcli -f inquire_config [ ] [-v vcenternodename] [-s tsmserverhostname] [-n vctrclinodename] [-p tsmserverport]
```

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

Parameters

[-v *vcenternodename*]

Depending on the backup environment, specify the virtual node 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 IBM Storage Protect server. If this parameter is not specified, the value in the profile is used.

[-n *vctrclinodename*]

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

[-p *tmsserverport*]

Specify the port of the IBM Storage Protect 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.

vSphere environment 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 IBM Storage Protect node configuration. As a result, the configuration is correct:

```
#TASK 38 inquire_config 20140108213337381
#PARAM INSTALLED=TSM
#RUN 32 20140108213337381
#LANG en_US
#PARAM BACKEND=TSM
#PARAM OPERATION_TYPE 5
#PHASE_COUNT 4
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=FVTSERIES11ESX6.STORAGE.MYCOMPANY.COM
#PARAM TSM_SERVER_PORT=1500
#PARAM TSMCLI_NODE_NAME=DPM02_VMCLI
#PARAM VCENTER_NODE_NAME=DPM02_VC1
#PARAM DATACENTER_NODE_NAME=
#PARAM OFFLOAD_HOST_NAME=
#PARAM PASSWORD_TYPE=CLINODE
#PARAM TSM_OPTFILE=C:\Users\ADMINI~1\AppData\Local\Temp\2\T4VBE42.tmp
#PARAM INPUT_FILE=
#PARAM TRACEFILE=
#PARAM TRACEFLAGS=
#PARAM RUNID=38
#PHASE INITIALIZE
#PHASE INQUIRE_DATACENTER_NODES
#CHILD datacenternode:DC1::DPM02_DC1
#PARENT vcenternode:DPM02_VC1
#PHASE INQUIRE_PROXY_NODES
#CHILD targetnode:DPM02_DC1
#PARENT peernode:DPM02_DC1_DM
#CHILD hladdress:tsmveesx2vm50.storage.mycompany.com
#PARENT peernode:DPM02_DC1_DM
#CHILD lladdress:49394
#PARENT peernode:DPM02_DC1_DM
#CHILD nodetype:DMNODE
#PARENT peernode:DPM02_DC1_DM
#CHILD partner:
#PARENT peernode:DPM02_DC1_DM
#CHILD targetnode:DPM02_DC1
#PARENT peernode:DPM02_DC1_2_MP_WIN
```

```
#CHILD hladdress:tsmveesx2vm50.storage.mycompany.com
#PARENT peernode:DPM02_DC1_2_MP_WIN
#CHILD lladdress:49453
#PARENT peernode:DPM02_DC1_2_MP_WIN
#CHILD nodetype:MPNODE
#PARENT peernode:DPM02_DC1_2_MP_WIN
#CHILD partner:DPM02_DC1_2_MP_LNX
#PARENT peernode:DPM02_DC1_2_MP_WIN
#CHILD targetnode:DPM02_DC1
#PARENT peernode:DPM02_DC1_2_MP_LNX
#CHILD hladdress:
#PARENT peernode:DPM02_DC1_2_MP_LNX
#CHILD lladdress:
#PARENT peernode:DPM02_DC1_2_MP_LNX
#CHILD nodetype:MPNODE
#PARENT peernode:DPM02_DC1_2_MP_LNX
#CHILD partner:DPM02_DC1_2_MP_WIN
#PARENT peernode:DPM02_DC1_2_MP_LNX
#PARAM STATUS=success
#PARAM STATUS=success
#END RUN 32 20140108213340100
#END TASK 38
#INFO FMM16014I The return code is 0.
#END
```

The PHASE INQUIRE_DATACENTER_NODES section shows the mapping of the datacenter name (DC1) from vSphere to the IBM Storage Protect node name for that datacenter (DPM02_DC1). The datacenter name is case sensitive and must be identical to the name shown in vSphere for the mapping to function.

The PHASE INQUIRE_PROXY_NODES section shows the data mover nodes with proxy access to each data center node. The format for this proxy relationship is shown in pairs:

```
#CHILD targetnode::<datacenter node name>
#PARENT peernode::<data mover node name>
```

Two types of proxy nodes are identified in the PHASE INQUIRE_PROXY_NODES section:

- The CHILD nodetype:DMNODE subsection identifies the data mover nodes and their proxy relationships.
- The CHILD nodetype:MPNODE subsection identifies the mount proxy nodes and their proxy relationships. These nodes represent the proxy system that accesses mounted VM disks through an iSCSI connection. Mount proxy nodes are required for file restore operations.

Inquire_detail

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

Syntax

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

```
vmcli -f inquire_detail -d datacenternodename | organizationvDCnodename [-a] [-n vmclinodename] [-o datamovernodename] [-p tsmserverport] [-e vmdetail] [-q dmverify | vmfs | vmsingle (-I inputfile)] [-s tsmserverhostname] [-vmrestoretype (instantrestore | instantaccess | alltype | mount)] [-w vmBackupLocation]
```

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

Parameters

-d *datacenternodename*

Specify the datacenter node name.

[-a]

Specify to show only the active backups on the IBM Storage Protect server.

[-n vmclinodename]

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

[-o datamovernodename]

Specify the data mover node name.

[-p tsmserverport]

Specify the port of the IBM Storage Protect 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.

[-e vmdetail]

Specify vmdetail to show more detailed information about the backed up VMs. For example, the parameter shows information about disks that are attached to the VM.

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

Specify to query the status of the data mover node identified by the **-o** parameter. You must specify the **-d** and **-o** parameters when you specify dmverify.

vmfs

Specify to query all VMware Virtual Machine File Systems (VMFS). This parameter shows high-level information about all VMs.

vmsingle

Specify to query individual VMs that are being restored during an instant access or instant restore operation.

-I inputfile

The *inputfile* value defines the full path and name of the input file. This keyword is valid with the **vmsingle** parameter only. Specify the name of the VM to query.

When the **q** option is not specified, the default value is vmfs. When the *inputfile* entry contains spaces, enclose the entry with quotation marks. For example:

```
-I "/my dir/my file"
```

[-s tsmserverhostname]

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

Windows**[-vmrestoretype (instantrestore | instantaccess | alltype | mount)]**

Specify this option to query active instant access or restore operations. This option also queries stale or orphan artifacts after a failure. The **vmrestoretype** parameter uses the following keywords:

instantrestore

The query lists VMs that are active in an instant restore operation.

instantaccess

The query lists VMs that are active in an instant access process.

alltype

The query lists VMs that are active in all instant operations.

mount

The query lists all active mount operations. For each mount operation, the output lists the mounted snapshots (restore points) that were created during a restore operation for a particular VM.

Restriction: When an instant restore or instant access operation that is issued from the backup-archive client (**dsmc**) is followed by an instant restore or instant access operation that is issued from the Data Protection for VMware command-line interface (**vmcli**) or Data Protection for VMware vSphere GUI, the recovery agent service must be restarted. This situation applies only when the **vmcli** accesses the server with a node name different from the one used by the data mover. This restriction applies to any order of operations between the two products.

Restart the service by going to **Start > Administrative Tools > Computer Management > Services and Applications > Services**. Look for service name IBM Storage Protect recovery agent in the **Services** window. The path to the **Services** window might vary depending on your operating system.

The service does not have to be restarted when the VMware datacenter name is specified with the `asnodename` option in the `dsm.opt` file.

[-w *vmBackupLocation*]

Specify the backup location or locations to query. You can specify one of the following values:

SERVER

The query is limited to backups that are on the IBM Storage Protect server. This value is the default.

LOCAL

The query is limited to persisted snapshots that are on the hardware storage.

BOTH

The query lists information for both backups that are on the IBM Storage Protect server and snapshots on the hardware storage.

This parameter is not valid when the **-vmrestoretype** parameter is also specified.

Example

In this example, the **vmcli -f inquire_detail** command is issued to query the VM named `antures` for details:

```
vmcli -f inquire_detail -s BORODIN.MAINZ.DE.IBM.COM -p 1505 -n JF_VMCLI_HANNE  
-v CHRISTO.MAINZ.DE.IBM.COM -o JF_MAINZ_DEVELOPMENT_DC_DM -d JF_MAINZ_DEVELOPMENT_DC  
-q vmsingle -I .\inputfile.txt --vmrestoretype (instantrestore | instantaccess)
```

The *inputfile* contains this statement:

```
vmname:antures
```

Set_domain

Use this **vmcli** command to apply changes to the domain settings.

Syntax

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

```
vmcli -f set_domain -I domainObjectListFile
```

Linux

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

The new domain value is stored in the **vmcli** database.

Parameters

-I *domainObjectListFile*

The *domainObjectListFile* has the following requirements:

- The file contains one VMware datacenter identifier per line.
- The valid identifier is the datacenter name.

If no domain is configured, the current instance is used to manage all datacenters that are available in the vCenter. When the `vmcli -f set_domain` command is run without `-I` parameter, the domain configuration is deleted.

An example *domainObjectListFile* is provided here:

```
#datacentername:datacenterName
datacentername:datacenterXYZ
datacentername:datacenterA*
datacentername:datacenterB*
...
```

Set_option

Use this **vmcli** command to set a parameter in the `vmcliprofile`.

Syntax

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

```
vmcli -f set_option [-m datacentermapping] [-n datamovernodename] [-p tsmserverport] [-s tsmserverhostname] [-v vctrnodename]
```

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

Parameters

-m datacentermapping

Specify the name of the data center that is associated with the datacenter node name (DC_name : : DC_nodename). The DC_name value is case sensitive and must match the name of your datacenter.

[-n datamovernodename]

Specify the data mover node name. This name is the node name for the IBM Storage Protect backup-archive client that is installed on the vStorage Backup Server. This node performs the data movement.

[-p tsmserverport]

Specify the port of the IBM Storage Protect 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.

[-s tsmserverhostname]

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

[-v vcenternodename]

Specify the vCenter node name. This node is the virtual node 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.

Example

In this example, the `vmcli -f set_option` command is issued to set the IBM Storage Protect server and its port:

```
vmcli -f set_option -s TEMPLE.MYCOMPANY.XYZ.COM -p 1650
```

The following output is displayed:

```
Setting VE_TSM_SERVER_NAME to: TEMPLE.MYCOMPANY.XYZ.COM
Setting VE_TSM_SERVER_PORT to: 1650
#INFO FMM16014I The return code is 0.
```

In this example, the `vmcli -f set_option` command is issued to set the data center mapping:

```
vmcli -f set_option -m DataCenter2::NANO_DATACENTER123
```

The following mapping is set in the profile:

```
VE_DATACENTER_NAME    DataCenter2::NANO_DATACENTER123
```

Set_password

Use this **vmcli** command to set the password for the guest VM.

Syntax

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

```
vmcli -f set_password [-type VMGuest] -I passwordfile
```

The **-type VMGuest** parameter is required when you set the password for application protection reporting.

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

Windows | **Linux** You must issue the **vmcli -f set_password** command before you run a guest scan operation.

Parameters

-type VMGuest

This parameter identifies that the password applies to a VM. This parameter is required when you set the password for application protection reporting.

-I *passwordfile*

Specify the following information in this file:

datacentername: *data center in vmcliprofile*

Specify the datacenter that contains the VM guests. The datacenter must be defined in the *vmcliprofile*. The password is applied to that datacenter only. For example:

```
datacentername:DataCenter1
```

username: *common VM guest user*

Specify the user name that logs in to the VM guest. For Windows, the *DOMAIN\User* format is allowed for the user name. For example:

```
username:Domain1\Administrator
```

password: *password*

Specify the password to log in to the VM guest.

The settings in the *passwordfile* must be specified on the same line.

Examples

Linux This example creates (or sets) a common VM guest name and password that is associated with DataCenter3. The *vmcliprofile* contains the following **VE_DATACENTER_NAME** settings:

```
VE_DATACENTER_NAME DataCenter1::TSM_DC1
VE_DATACENTER_NAME DataCenter2::TSM_DC2
```



```
VE_DATACENTER_NAME DataCenter3::TSM_DC3
VE_DATACENTER_NAME DataCenter4::TSM_DC4
```

The *passwordfile* contains the following settings. The settings in the *passwordfile* must be specified on the same line:

```
datacentername:DataCenter3 username:tdpvmwareuserY password:tdpvmwareuserYpwd
```

As a result, the **vmcli -f set_password -type VMGuest -I password.txt** sets the password as shown in the command output:

```
IBM Storage Protect Command Line Wrapper for Virtual Environments Version: 8.1.0
Build Date: Mon Dec 12 20:03:31 2016
IBM Storage Protect API Version 81000
IBM Storage Protect Command Line Wrapper Compile Version 81000
#PARAM OPERATION_TYPE 8
#PHASE_COUNT 3
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=ORION.FINANCE.MYCOMPANY.COM
#PARAM TSM_SERVER_PORT=1500
#PARAM TSMCLI_NODE_NAME=KA3095_TSMCLI_SLUDGE
#PARAM VCENTER_NODE_NAME=
#PARAM DATACENTER_NODE_NAME=
#PARAM OFFLOAD_HOST_NAME=
#PARAM TSM_OPTFILE=/tmp/T4VE_OD3PZ9
#PARAM INPUT_FILE=/opt/tivoli/tsm/tdpvmware/common/scripts/password.txt
#PARAM TRACEFILE=
#PARAM TRACEFLAGS=
#PHASE INITIALIZE
#PHASE SET_PASSWORD
STATUS=success
#END
```

Windows This example creates (or sets) a common VM guest name and password that is associated with DataCenter1. The *vmcli* profile contains the following **VE_DATACENTER_NAME** settings:

```
VE_DATACENTER_NAME DataCenter1::TSM_DC1
VE_DATACENTER_NAME DataCenter2::TSM_DC2
```

The *passwordfile* contains the following settings. The settings in the *passwordfile* must be specified on the same line:

```
datacentername:DataCenter1 username:Domain1\Administrator password:secret1
```

As a result, the **vmcli -f set_password -type VMGuest -I password.txt** sets the password as shown in the command output:

```
IBM Storage Protect Command Line Wrapper for Virtual Environments Version: 8.1.0
Build Date: Mon Dec 12 20:03:31 2016
IBM Storage Protect API Version 81000
IBM Storage Protect Command Line Wrapper Compile Version 81000
#PARAM OPERATION_TYPE 8
#PHASE_COUNT 3
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=ORION.FINANCE.MYCOMPANY.COM
#PARAM TSM_SERVER_PORT=1500
#PARAM TSMCLI_NODE_NAME=KA3095_TSMCLI_SLUDGE
#PARAM VCENTER_NODE_NAME=
#PARAM DATACENTER_NODE_NAME=
#PARAM OFFLOAD_HOST_NAME=
#PARAM TSM_OPTFILE=/tmp/T4VE_OD3PZ9
#PARAM INPUT_FILE=C:\Program Files\IBM\StorageProtect\Framework\VEGUI
scripts\password.txt
#PARAM TRACEFILE=
#PARAM TRACEFLAGS=
#PHASE INITIALIZE
#PHASE SET_PASSWORD
STATUS=success
#END
```

Windows When you create the password file by 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*) by specifying the **echo** command:

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

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

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

Windows | **Linux** This example sets the password in file *pwd.txt* for domain *mydomain* and user *user1*:

```
set -f set_password -I pwd.txt -pwtype domain -domain mydomain -user user1
```

Get_password_info

Use this **vmcli** command to view the status of guest credentials that are set for the managed datacenters.

Syntax

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

```
vmcli -f get_password_info -type VMGuest
```

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

Parameters

-type VMGuest

This required parameter identifies that the password information is returned for a guest VM. The **username** value (shown in the **#CHILD** statement) of the command output confirms that the password is set for that **username**. The **datacentername** value (shown in the **#PARENT** statement) of the command output identifies the associated datacenter for which the password is set.

Example

Windows This example shows the status of the managed datacenters that are associated with the VM guest:

```
vmcli -f get_password_info -type VMGuest
```

```
#TASK 0 get_password_info 20130129162344670
#RUN 0 20130129162344685
#LANG en_US
#PARAM BACKEND=TSM
#PARAM OPERATION_TYPE 4
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=
```

```
#PARAM TSM_SERVER_PORT=
#PARAM TSMCLI_NODE_NAME=
#PARAM VCENTER_NODE_NAME=
#PARAM DATACENTER_NODE_NAME=
#PARAM OFFLOAD_HOST_NAME=
#PARAM PASSWORD_TYPE=VMGUEST
#PARAM TSM_OPTFILE=C:\Users\ADMINI~1\AppData\Local\Temp\2\T4V3B15.tmp
#PARAM INPUT_FILE=
#PARAM TRACEFILE=
#PARAM TRACEFLAGS=
#CHILD username:<mydomain\myuser>
#PARENT datacentername:DataCenter1
#CHILD username:<mydomain\myuser>
#PARENT datacentername:DataCenter2
#CHILD username:<mydomain\myuser>
#PARENT datacentername:DataCenter3

#PARAM STATUS=success
#END RUN 0 2013012916234513
#END TASK 0
#INFO FMM16014I The return code is 0.
#END
```

Start_guest_scan

Use this **vmcli** command to scan guest VMs for application information.

The **vmcli -f start_guest_scan** command saves VM name, application, and globally unique identifier (GUID) information to the IBM Storage Protect server.

You must issue the **vmcli -f set_password** command before you run a guest scan operation.

Syntax

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

```
vmcli -f start_guest_scan -dcscan "datacenterNvmcliprofile,...," | ALL_DC -o
datamovernodename
```

Required Parameters

-dcscan datacenterNvmcliprofile | ALL_DC

Specify one or more datacenter names that are defined in the *vmcliprofile*. Repeat datacenter names with a comma. Double quotation marks (") must be specified at the beginning and at the end of the datacenter name list. For example:

```
-dcscan "Local DC,svc"
```

To scan all VM guests in all datacenters, specify the **ALL_DC** parameter.

-o datamovernodename

Specify the data mover node that is configured with proxy authority access to the datacenters specified by **-dcscan**.

During a **vmcli -f start_guest_scan** operation, Data Protection for VMware copies files to a temporary subdirectory in the remote directory (\$TEMP_REMOTE\TSMSCAN) on the guest VM. The remote directory must be unlocked and not used by another application. Data Protection for VMware determines the location of the remote directory in the following order:

1. If the **TEMP** environment variable is set, **TEMP_REMOTE** is set as the **TEMP** environment variable.
2. If the **TEMP** environment variable is not set, **TEMP_REMOTE** is set as C:\TEMP.

Example

Windows In this example, the *vmcliprofile* contains the following **VE_DATACENTER_NAME** settings:

```
VE_DATACENTER_NAME: DataCenter1:TSM_DC1
VE_DATACENTER_NAME: DataCenter2:TSM_DC2
```

The data mover node, VC1_DC1_DM1, is configured with proxy authority access to DataCenter1 and DataCenter2.

Windows The following command is issued to scan all guest VMs in DataCenter1 and DataCenter2:

```
vmcli -f start_guest_scan -dcscan "DataCenter1,DataCenter2" -o VC1_DC1_DM1
```

The following application information is displayed:

```
IBM Storage Protect Command Line Wrapper for Virtual Environments
Version: 8.1.0
Build Date: Mon Dec 12 20:03:31 2016
IBM Storage Protect API Version 81000
IBM Storage Protect Command Line Wrapper Compile Version 81000
#PARAM OPERATION_TYPE 9
#PHASE_COUNT 4
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=OREO.STORE.XYZ.COM
#PARAM TSM_SERVER_PORT=1500
#PARAM TSMCLI_NODE_NAME=VC1_VCLI1
#PARAM VCENTER_NODE_NAME=VC1
#PARAM DATACENTER_NODE_NAME=VC1_DC1
#PARAM OFFLOAD_HOST_NAME=VC1_DC1_DM1
#PARAM PASSWORD_TYPE=CLINODE
#PARAM TSM_OPTFILE=C:\Users\ADMINI~1\AppData\Local\Temp\2\T4V9393.tmp
#PARAM INPUT_FILE=
#PARAM TRACEFILE=c:\amd64_unicode\tsmcli.trace
#PARAM TRACEFLAGS=service,VMVCB,VTSMVSS,verbdetail,C2C
#PHASE INITIALIZE
#CHILD targetnode:VC1_DC1
#PARENT peernode:VC1_DC1_DM1
#CHILD hladdress:9.52.62.65
#PARENT peernode:VC1_DC1_DM1
#CHILD lladdress:50408
#PARENT peernode:VC1_DC1_DM1
#PHASE READ_DATACENTER_GUEST_PASSWORD
#PHASE SCANGUEST
#PARAM STATUS=success
```

```
#CHILD scanid: DataCenter1::VC1_DC1.1358316054281
#PARENT datacentername: DataCenter1::VC1_DC1
#PARAM OPERATION_TYPE 9 #PHASE_COUNT 4
#PHASE PREPARE
#PARAM BACKUP_TYPE=0
#PARAM TSM_SERVER_NAME=OREO.STORE.XYZ.COM
#PARAM TSM_SERVER_PORT=1500
#PARAM TSMCLI_NODE_NAME=VC1_VCLI1
#PARAM VCENTER_NODE_NAME=VC1
#PARAM DATACENTER_NODE_NAME=VC1_DC2
#PARAM OFFLOAD_HOST_NAME=VC1_DC1_DM1
#PARAM PASSWORD_TYPE=CLINODE
#PARAM TSM_OPTFILE=C:\Users\ADMINI~1\AppData\Local\Temp\2\T4V50B.tmp
#PARAM INPUT_FILE= #PARAM TRACEFILE=c:\amd64_unicode\tsmcli.trace
#PARAM TRACEFLAGS=service,VMVCB,VTSMVSS,verbdetail,C2C
#PHASE INITIALIZE #CHILD targetnode:VC1_DC2
#PARENT peernode:VC1_DC1_DM1
#CHILD hladdress:9.52.62.65
#PARENT peernode:VC1_DC1_DM1
#CHILD lladdress:50408
#PARENT peernode:VC1_DC1_DM1
#PHASE READ_DATACENTER_GUEST_PASSWORD
#PHASE SCANGUEST
#PARAM STATUS=success
#CHILD scanid: DataCenter2::VC1_DC2.1358316054281
#PARENT datacentername:DataCenter2::VC1_DC2
#INFO FMM16014I The return code is 0.
#END
```

Windows The #PARAM STATUS=success message (in the #PHASE SCANGUEST section) confirms only that the datacenter was successfully submitted for processing by the data mover. The actual scan status for each VM is available only after the data mover completed processing that VM. To view the scan status of an individual VM, see the value in the Scan Status column of the Data Protection for VMware vSphere

GUI **Application Configuration Status** report. To view the overall status of the scan operation, see the **Overall Scan Status** value in the Data Protection for VMware vSphere GUI **Report** window.

Important: If you receive an error after you run the **vmcli -f start_guest_scan** command, view the contents of the `dsmererror.log` file for more information. The `dsmererror.log` file is on the system that is associated with the data mover node defined by the **OFFLOAD_HOST_NAME** parameter in the command output. By default, error log files are in the installation directory:

C:\Program Files\Tivoli\TSM\baclient

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 vSphere GUI is installed:

Linux /home/tdpvmware/tdpvmware/config

Windows C:\Program Files\IBM\StorageProtect\Framework\VEGUI\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 /home/tdpvmware/tdpvmware

Windows DERBY_HOME C:\Program Files\IBM\StorageProtect\Framework\VEGUI\derby

VE_DATACENTER_NAME <data_center_name::DATA_CENTER_NODE_NAME>

Specify the VMware datacenter (`datacenter name`) with a value that is case-sensitive and that matches the datacenter name used in the vCenter. Specify the virtual node (`DATA_CENTER_NODE_NAME`) that maps to the datacenter. If the vCenter manages several datacenters, you can specify this parameter for each datacenter.

Example:

```
VE_DATACENTER_NAME  DataCenter1::Fin_Datacenter1
VE_DATACENTER_NAME  DataCenter2::Fin_Datacenter2
VE_DATACENTER_NAME  DataCenter3::Fin_Datacenter3
```

Important: The Data Protection for VMware vSphere GUI does not support datacenters with the same name in the vCenter.

This parameter is valid only in a vSphere environment.

Restrictions:

Data Protection for VMware does not support a one-to-many or many-to-one relationship between the datacenter name and the datacenter node. For example, the following relationships are not supported:

```
VE_DATACENTER_NAME  DataCenter1::Fin_Datacenter1
VE_DATACENTER_NAME  DataCenter1::Fin_Datacenter2
VE_DATACENTER_NAME  DataCenter1::Fin_Datacenter3
```

Or

```
VE_DATACENTER_NAME  DataCenter1::Fin_Datacenter1
VE_DATACENTER_NAME  DataCenter2::Fin_Datacenter1
VE_DATACENTER_NAME  DataCenter3::Fin_Datacenter1
```

Data Protection for VMware support for VM backup and restore operations is limited to VM names and datacenter names that contain English 7-bit ASCII characters only. VM names and datacenter names that use other language characters are not currently supported. Additional character restrictions are listed in [Appendix A, “Troubleshooting,”](#) on page 163.

After a datacenter name is created and associated with the IBM Storage Protect node, be aware of these restrictions:

- Do not change the datacenter name in the vCenter without also creating the IBM Storage Protect node name and associating it with the new datacenter name.
- Do not change the datacenter name and the profile without also changing the IBM Storage Protect node name.
- Do not create a datacenter mapping value in the profile with a previously used IBM Storage Protect node.

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

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

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

Example:

```
VE_TSMCLI_NODE_NAME  VC1_VCLI1
```

Restriction: The VMCLI node does not support the SSL protocol or LDAP authentication when communicating with the IBM Storage Protect server.

VE_TSM_SERVER_NAME <server host name or IP address>

Specify the host name or IP address of the IBM Storage Protect 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 IBM Storage Protect server. The default value is 1500.

Example:

```
VE_TSM_SERVER_PORT 1500
```

VE_TSM_SSL YES|NO

Specify whether to enable Secure Sockets Layer (SSL) to provide secure client and server communications. No is the default.

VE_TSM_SSLACCEPTCERTFROMSERV YES|NO

Specify whether the backup-archive client or the API application accept and trust the IBM Storage Protect server's Secure Sockets Layer (SSL) public certificate the first time they connect. YES is the default. You can use this option to connect only to an IBM Storage Protect server version 8.1.2 and later.

VE_TSM_SSLREQUIRED DEFAULT|YES|NO|SERVERONLY

Specify the conditions when SSL is or is not required when the client logs on to the IBM Storage Protect server or storage agents. DEFAULT is the default. When communicating with the IBM Storage Protect server version 8.1.2 and later, this option no longer applies since SSL is always used.

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 the IBM Storage Protect server, 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 to begin. 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.

Example Linux profile in a vSphere environment

Linux

```
VE_TSM_SERVER_NAME      9.11.90.28
VE_TSM_SERVER_PORT      1500
VE_TSMCLI_NODE_NAME     my_vc1_vcli1
VE_VCENTER_NODE_NAME    my_vc1
VE_DATACENTER_NAME      Clovis Lab::MY_VC1_DC1
VMCLI_TASK_EXPIRATION_TIME 864000 # in seconds, defaults to 864000s = 10 days
VMCLI_RESTORE_TASK_EXPIRATION_TIME 2592000 # in seconds, defaults to 2592000s = 30 days
VMCLI_GRACE_PERIOD      2592000 # in seconds, defaults to 2592000s = 30 days
VMCLI_SCHEDULER_INTERVAL 60 # in seconds, defaults to 1s
VMCLI_DB_HOST            localhost
VMCLI_DB_PORT            1527
VMCLI_CACHE_EXPIRATION_TIME 600 # in seconds, defaults to 600s = 10 min
VMCLI_DB_NAME            VMCLIDB
VMCLI_RECON_INTERVAL_FCM 600 # setting in seconds default 600s = 10 min
VMCLI_RECON_INTERVAL_TSM 1200 # setting in seconds default 1200s = 20 min
VMCLI_DB_BACKUP          AT 00:00
VMCLI_DB_BACKUP_VERSIONS 3
VMCLI_LOG_DIR            logs
DERBY_HOME               /home/tdpvmware/tdpvmware
```

Example Windows profile in a vSphere environment

Windows

```

VE_TSM_SERVER_NAME      philadelphia      # -s
VE_TSM_SERVER_PORT      1500              # -p
VE_TSMCLI_NODE_NAME     CLI_WIN8x32        # -n
VE_VCENTER_NODE_NAME    VC_WIN8x32         # -v
VE_DATACENTER_NAME      DC_CVT::DC_Win8x32
VMCLI_TASK_EXPIRATION_TIME 864000 # in seconds, defaults to 864000s = 10 days
VMCLI_RESTORE_TASK_EXPIRATION_TIME 2592000 # in seconds, defaults to 2592000s = 30 days
VMCLI_GRACE_PERIOD       2592000 # in seconds, defaults to 2592000s = 30 days
VMCLI_SCHEDULER_INTERVAL 60 # in seconds, defaults to 1s
VMCLI_DB_HOST            localhost
VMCLI_DB_PORT            1527
VMCLI_CACHE_EXPIRATION_TIME 600 # in seconds, defaults to 600s = 10 min
VMCLI_DB_NAME            VMCLIDB
VMCLI_RECON_INTERVAL_FCM 600 # setting in seconds default 600s = 10 min
VMCLI_RECON_INTERVAL_TSM 1200 # setting in seconds default 1200s = 20 min
VMCLI_DB_BACKUP          AT 00:00
VMCLI_DB_BACKUP_VERSIONS 3
VMCLI_LOG_DIR            logs
DERBY_HOME               C:\Program Files\IBM\StorageProtect\Framework\VEGUI\derby

```

Recovery Agent command-line interface

Use the Recovery Agent command-line interface (CLI) to access Data Protection for VMware functions.

The Recovery Agent CLI can be viewed as a command-line API to the IBM Storage Protect recovery agent. Changes completed with the Recovery Agent CLI to the recovery agent take effect immediately.

You can use the Recovery Agent CLI to manage only one system running the recovery agent.

Starting the Recovery Agent command-line interface

Start the Recovery Agent CLI from the Windows **Start** menu.

About this task

To start the Recovery Agent CLI, complete the following steps:

Procedure

1. From the Windows **Start** menu, click **Programs > IBM Storage Protect > Data Protection for VMware > IBM Storage Protect recovery agent**.
2. In the command prompt window, enter one of the following commands:
 - To run the Recovery Agent CLI:

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

- **Windows** To display the help for the Recovery Agent CLI:

```
RecoveryAgentShell.exe -h
```

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

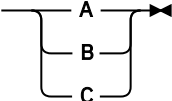
Variables

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 —  ►►

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:



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 command

Use the **mount** command to complete various 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, the IBM Storage Protect recovery agent must be running. Use the **set_connection** command to connect a RecoveryAgentShell.exe to the mount application.

Snapshots are mounted or unmounted on the system where the 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 the 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>"
```

When you use the recovery agent in an iSCSI network, and the Recovery Agent does not use a data mover, go to the C:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\RecoveryAgent.conf file and specify the [IMOUNT] tag and **Target IP** parameter:

```
[IMOUNT config]
Target IP=<IP address of the network card on the system
that exposes the iSCSI targets.>
```

For example:

```
[General config]
param1
param2
...
[IMount config]
Target IP=9.11.153.39
```

After you add or change the Target IP parameter, restart the Recovery Agent GUI or Recovery Agent CLI.

- **-rep** - This tag is required.

Use it to specify the IBM Storage Protect server that is storing the VMware snapshots, and the IBM Storage Protect node that has access to the VMware backups. For example:

```
tsm: ip=<ip/host_name> port=<port_number>
node=<node_name> pass=<node_password>
```

You can also specify the `as_node` and `from_node` options. If the `password` field is empty, the recovery agent attempts to use the password for the stored node.

- **-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 it to specify the date of the snapshot that you want to mount. The date format is `yyyy-Mmm-dd hh:mm:ss`. For example:

```
-date "2013-Apr-12 22:42:52 AM"
```

To view the active (or latest) snapshot, specify `last` snapshot.

- **-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**).
- **-disk** - This tag is required. Use it to specify the disk number of the source backed up VMware machine to be mounted.
- **-ExpireProtect** - This tag is optional. During a mount operation, the snapshot on the IBM Storage Protect server is locked to prevent it from expiring during the operation. Expiration might occur because another snapshot is added to the mounted snapshot sequence. This value specifies whether to disable expiration protection during the mount operation. You can specify one of the following values:

Yes

Specify Yes to protect the snapshot from expiration. This value is the default. The snapshot on the server is locked and the snapshot is protected from expiration during the mount operation.

No

Specify No to disable expiration protection. The snapshot on the server is not locked and the snapshot is not protected from expiration during the mount operation. As a result, the snapshot might expire during the mount operation. This expiration can produce unexpected results and negatively impact the mount point. For example, the mount point can become unusable or contain errors. However, expiration does not affect the current active copy. The active copy cannot expire during an operation.

When the snapshot is on a target replication server, the snapshot cannot be locked because it is in read-only mode. A lock attempt by the server causes the mount operation to fail. To avoid the lock attempt and prevent such a failure, disable expiration protection by specifying No.

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 "2014-Jan-21 10:46:57 AM -ExpireProtect=Yes"
```

In this example, a snapshot of VMware named VM-03ent is located on a server with IP 10.10.10.01. Disk number 1 of this snapshot is mounted to the system where the recovery agent is running.

del

Use this command type to dismount one or all mounted backups from the system where the 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, repare point, or iSCSI target created using the **mount** command. Use the *everything* variable 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 target, 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 IBM Storage Protect server storing the VMware snapshots, and to specify the IBM Storage Protect 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 VMs.

```
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 the IBM Storage Protect server. A connection cannot be removed when it is in use, such as when mounted volumes exist. There is only one tag for the **remove** type:

-rep - This tag is required. Use this tag to specify the IBM Storage Protect server connection to be removed.

In the following example, remove the connection to a server (10.10.10.01) using node NodeName:

```
mount remove -rep "tsm: NodeName@ip"
```

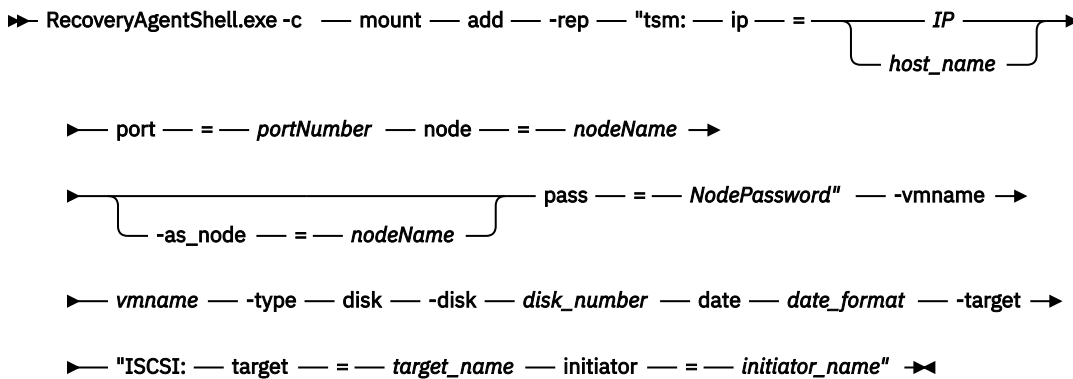
view

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
```

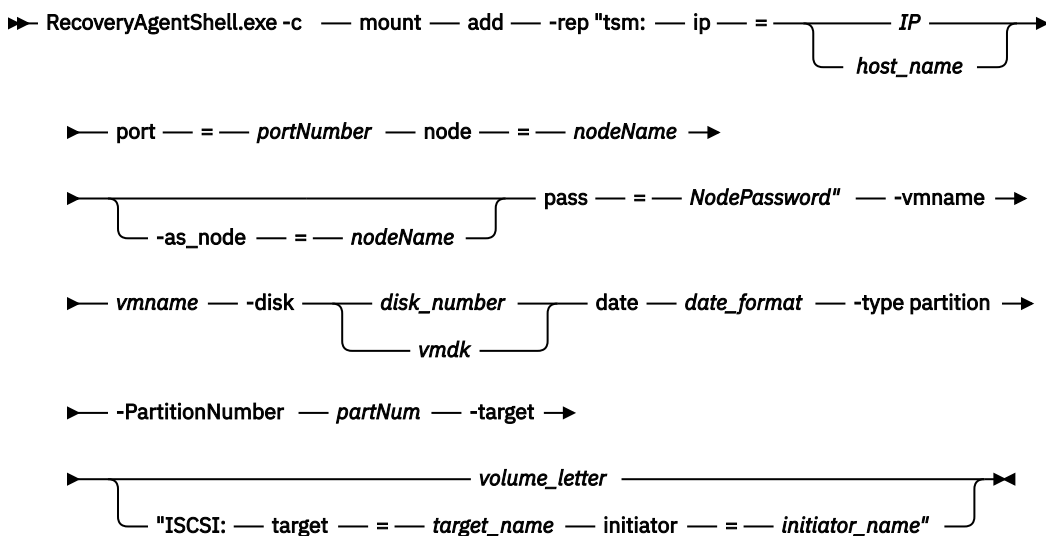
Mounting a disk

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 command

The **set_connection** command sets the Recovery Agent CLI to work with a specified recovery agent. Use the following format for the **set_connection** command:

```
RecoveryAgentShell.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 recovery agent connection.

In the following example, the Recovery Agent CLI is set to work with recovery agent on the *ComputerName* host.

```
set_connection mount_computer ComputerName
```

Help command

The **help** command displays the help for all of the supported Recovery Agent CLI commands.

Use the following format for the **help** command:

```
RecoveryAgentShell.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 9. 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 dismounted 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.
1	FBC_MSG_MOUNT_FAIL	Mount failed (See the mount logs for details).
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.

Table 9. Recovery Agent CLI return codes (continued)

Return Code	Value	Description
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 IBM Storage Protect 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.
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. For more information, see the documentation.

Chapter 6. Backing up VMware data

Use Data Protection for VMware to store backup versions of your VMware virtual machines on the IBM Storage Protect server.

Restriction: The following restrictions apply to VMware VMDKs that are involved in a backup operation:

- For incremental forever backup mode, each individual VMDK involved in a backup operation cannot exceed 8 TB. If a VMDK exceeds 8 TB, the backup operation fails. To increase the size of the VMDK to be larger than the default 2 TB, specify the maximum size with the `vmmaxvirtualdisks` option. For more information, see [Vmmaxvirtualdisks](#).

To prevent a failure during either backup mode, you can skip processing the VMDK by specifying `vmskipmaxvirtualdisks yes` in the data mover options file. For more information, see [Vmskipmaxvirtualdisks](#).

Backing up virtual machine data to IBM Storage Protect

Create a run now backup task or scheduled backup task for your virtual machine data. The data is stored on IBM Storage Protect server storage.

Before you begin

Note the following usage tips.

Note: Before you back up virtual machines that are hosting Microsoft Exchange Server databases, mount the Exchange databases.

Note: During snapshot creation, vCenter may be unresponsive for a short time. This can cause communication failure since all communication is done via vCenter by default. The best ways to avoid this are:

- Connect directly to the ESXi host that hosts vCenter. For example:

```
dsmc backup VM VcenterVM1 -mode=ifull -vmchost=<esxi ip>
```

or

```
fqdn> -vmcuser=root -vmcpw=<pwd of root>
```

- Reduce the load on vCenter so that a snapshot is created fast enough that vCenter responds before timeout.
- Increase the resources allocated to vCenter to speed up the operation and avoid timeout.

About this task

- During backup processing, Data Protection for VMware bypasses a guest Exchange Server database that is dismounted, corrupted, or in a Suspend state in a Database Availability Group (DAG). Databases in these states are excluded from virtual machine backups and are not available to restore.
- A run now backup task can be issued immediately or at a later time. A scheduled backup task is issued by the IBM Storage Protect server at the scheduled time. It cannot be issued immediately.

Procedure

Back up your virtual machine by following these steps:

1. Start the Data Protection for VMware vSphere GUI with either of these methods:
 - Click the Data Protection for VMware vSphere GUI icon in the **Solutions and Applications** window of the vSphere Client.

- Open a web browser, and go to the GUI web server. For example:

```
https://guihost.mycompany.com:9081/TsmVMwareUI/
```

Log on with the vCenter user ID and password.

2. In the **Getting Started** window, click **Define a backup task** to open the **Schedule a Backup** wizard. Review the welcome information and click **Next** to create the backup task.
3. Click **Next** to begin the wizard. Follow the instructions in each page of the wizard and click **Next** to continue to the next page.
4. In the **General** page, specify a name for the backup schedule you are creating in the **Backup schedule name** field. The schedule name must not contain any spaces. You can optionally add a description for the schedule.
5. In the **Source** page, select a VMware datacenter and expand the branches in the navigation tree. Select the virtual machines, virtual machine templates, host cluster, or host that you want to back up. 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 rename a virtual machine, and you select this option, the schedule backs up the renamed virtual machine.
 - If you do not select this option, only virtual machines can be selected for backup. Host clusters and hosts cannot be selected.
 - To easily exclude virtual machines that follow a naming convention from being added, expand the **Advanced VM filter option**. Enter a text pattern that identifies the virtual machines to exclude. Specify an asterisk (*) to match any character. Specify a question mark (?) to match a single character. For example:

```
vm=prod1*,*testvm,*dept*, dept4?prod
```

Click **Apply filter** to disable selection for these virtual machines. If you enter a virtual machine name without wildcard characters, and the virtual machine is known in the source tree, then this virtual machine is removed from the filter display. However, it is shown in the tree as not checked.

6. In the **Destination** page, select the data mover node that runs the backup operation. Select a data mover node that is not currently used in a backup or restore operation.
7. In the **Schedule** page, specify when to run the backup by clicking **Run the backup now** or **Schedule for later**. If your user authority is insufficient, **Schedule for later** is unavailable. Select the appropriate **Backup strategy** from the drop-down list:

IFINCREMENTAL / IFFULL

Indicates the incremental-forever incremental backup type and the incremental-forever full backup type. These backup types are applicable only if you have a license to use IBM Storage Protect for Virtual Environments.

Select the appropriate **Backup type**:

- Click **Incremental** to back up the blocks that changed since the previous backup (full or incremental). The most recent incremental is appended to the previous backup. If a full backup does not exist for this VM, a full backup is automatically performed. As a result, you do not have to verify that a full backup exists.

- Click **Full** to create an image of an entire VM. After the full backup is taken, there is no requirement to schedule additional full backups. When full is selected, VM templates that are unchanged since the last backup are also included.
 - If you selected **Schedule for later**, click **Next** and proceed to Step 8.
 - If you selected **Run the backup now**, click **Next**, and proceed to Step 9.
8. In the **Repetition** page, specify the following information:
 - a) If you selected **A full backup, followed by six incremental backups** in the previous step:
 - i) Specify the date and time to run the first backup. The first full backup is scheduled to run at this date and time.
 - ii) The six incremental backups are scheduled to run on the remaining six days of the week and at the selected time.
 - b) If you selected **Incremental** or **Full** in the previous step:
 - i) Specify the date and time to run the first backup.
 - ii) Specify the interval that you want the backup to run.
 - c) Click **Next**.
 9. In the **Summary** page, review your backup settings and click **Finish** to save your task. If you selected **Run the backup now**, the backup operation begins immediately when you click **Finish**.

What to do next

After the backup schedule has completed, you can verify that the virtual machines have been backed up in the **Reports** tab.

Setting options for an incremental forever backup schedule

When you schedule incremental forever backups, you can ensure that there are frequent backups of the VMs and reduce the size of each backup.

Before you begin

Ensure that client-side data deduplication is enabled for the storage pool.

Procedure

1. Start a data mover command-line session:
 - **Windows** Open a command prompt and change to the data mover installation directory. For example:


```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```
 - **Linux** Open a terminal window and change to the data mover installation directory. For example:


```
cd /opt/tivoli/tsm/client/ba/bin
```
2. Edit the IBM Storage Protect client backup-archive client options file.
 - **Windows** Specify these options in the `dsm.opt` options file.
 - **Linux** Specify these options in the `dsm.sys` file in the stanza for the data mover node.
 - a) Enable compression by adding the option `compression yes` to the file.
 - b) Enable deduplication by adding the option `deduplication yes` to the file.
 - c) Modify the trigger values for megablock refreshes by setting one of the following options:
 - Enable a number of objects as the trigger by adding `mbobjrefreshthresh number` to the file.

- Enable a percentage of objects as the trigger by adding `mbpctrefreshtresh percentage` to the file.

For more information, see the data mover `mbobjrefreshtresh` and `mbpctrefreshtresh` options in [Client options reference](#).

3. Repeat Step 2 for each VMware guest.

Backing up virtual machine data to local disks

You can create a schedule to back up your VM data and store it on a storage device that is configured in the IBM Storage Protect Snapshot for VMware profile. You can run the schedule immediately or at a later time.

Before you begin

Before you back up virtual machines that are hosting Microsoft Exchange Server databases, mount the Exchange databases.

About this task

During backup processing, Data Protection for VMware bypasses a guest Exchange Server database that is dismounted, corrupted, or in a Suspend state in a Database Availability Group (DAG). Databases in these states are excluded from virtual machine backups and are not available to restore.

Procedure

Create a schedule to back up VM by following these steps:

1. Start the Tivoli® Storage FlashCopy® Manager for VMware GUI with either of these methods:

- Click the Tivoli Storage FlashCopy Manager for VMware GUI icon in the **Solutions and Applications** window of the vSphere Client.
- Open a web browser, and go to the GUI web server. For example:

```
https://guihost.mycompany.com:9081/TsmVMwareUI/
```

Log on with the vCenter user ID and password.

2. In the **Getting Started** window, click **Define a backup task** to open the **Schedule a Backup** wizard. Review the welcome information and click **Next** to create the backup task.
3. In the **Backup settings** page, select **Back up to IBM Storage Protect Snapshot** as the **Type of backup to run**.
4. In the **General** page, specify a name for the backup schedule you are creating in the **Backup Schedule Name** field. The schedule name must not contain any spaces. You can optionally add a description for the schedule.
5. In the **Source** page, select a VMware datacenter to view all VMs available for that datacenter. You can also expand the datacenter and select a VMware datastore to view all VMs available for that datastore. Select the VMs or VM templates to back up in the selection table. If you do not know the location of a VM, select a datacenter to view all its VMs.

Tip: Fine-tune the list of VMs by entering text in the **Filter** field. The wildcard character (*) is not allowed.

6. In the **Settings** page, select the type of VMware snapshot to take when backing up to the local IBM Storage Protect Snapshot for VMware disk.
7. In the **Destination** page, select the device class on the local disk of IBM Storage Protect Snapshot for VMware to back up to. The device class defines the storage device and target disks that are used to create the hardware snapshot of the logical units (LUNs) that are used for the datastores. The device classes must be defined in the IBM Storage Protect Snapshot for VMware configuration before they are available in the list.

8. In the **Schedule** page, you can select to run the backup immediately or schedule the backup to run at a later time.
- To begin the backup at the completion of this wizard, select **Execute immediately**.
 - To schedule your backup to run at a later time, do the following steps:
 - a. Select **Schedule for later (repeated) execution**.
 - b. In the **Date and time of the first backup** field, specify when you want to run the schedule for the first time.
 - c. Specify the interval that you want the schedule to run in the **How often to run the backup** field.
 - d. In the **Date and time of the last backup** field, specify when you want to stop the schedule from running. If you do not specify the date and time of the last backup, the schedule creation cannot continue.
 - e. If you want to include newly added or discovered VMs to future runs of this schedule, select **Newly added virtual machines are included in this backup task**. The table lists the datastores that contain the newly added VMs. This check box has the following characteristics when selected:
 - If you select all of the VMs on one datastore or datacenter, the schedule backs up that datastore or datacenter. That is, all VMs on that datastore or datacenter, present and future, are backed up.
 - If you select one or more VMs on a datastore or datacenter (but not all), then the schedule includes only the selected VMs and any future VMs that are added to the datastore or datacenter. The remaining VMs on that datastore or datacenter that are not selected are excluded.
 - If you rename a VM, and you select this option, the schedule backs up the renamed VM.
- To modify the list of datastores, go to the **Source** page and select VMs from only the preferred datastores. Clear the selected VMs from the non-preferred datastores.
9. In the **Summary** page, review your backup settings and click **Finish** to save your schedule. If you selected to run the backup now, the backup operation begins immediately when you click **Finish**.

What to do next

After the backup schedule completes, you can verify that the VMs were backed up in the **Reports** tab.

Backing up migrated virtual machines

When you migrate virtual machines to a different VMware datacenter or vCenter server, you can back up the virtual machines.

Before you begin

To back up migrated virtual machines, meet the following prerequisites:

- The migrated virtual machine must be running in a VMware vSphere 6 environment.
- Before you migrate the virtual machine, back up the virtual machine with Data Protection for VMware V7.1.3 or later. Verify that the backup completed without error.
- On the virtual machine, verify that VMware Storage vMotion is installed.

The following environment migrations are supported:

- Migrate from one vCenter to another vCenter. For example: vCenter A, datacenter A, and data mover A migrated to vCenter B, datacenter B, and data mover B.
- Migrate from one datacenter to another datacenter within the same vCenter. For example: vCenter A, datacenter A, and data mover A migrated to vCenter A, datacenter C, and data mover C.

About this task

To migrate a virtual machine from one VMware datacenter to another datacenter within the same vCenter server, complete the following steps:

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\tsm\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Create a full VM backup of the migrated virtual machine.

Store the backup on the original datacenter node from where the virtual machine was migrated.

For example, if virtual machine VM_1 was migrated from VMware datacenter DC_A to datacenter DC_C, then back up VM_1 from data mover DM_A to datacenter DC_C directly. The following sample command is provided:

```
dsmc backup vm VM_1 -vmbackuptype=fullvm -nodename=DC_C nodename  
-password=DC_C nodename_password
```

3. Deactivate the active backup of the migrated virtual machine on the original datacenter.

Issue this command on the original datacenter node from where the virtual machine was migrated.

For example, if virtual machine VM_1 was migrated from VMware datacenter DC_A to datacenter DC_C, then issue this command on datacenter node DC_A. The following sample command is provided:

```
dsmc expire -objtype=vm VM_1 -nodename=DC_A nodename  
-password=DC_A nodename_password
```

Backing up organization vDCs to IBM Storage Protect

You can create an immediate backup (**Back Up Now**) or schedule a backup (**Create Backup Schedule**) for an organization vDC. The data that is backed up is stored on IBM Storage Protect server storage.

Procedure

Back up your organization vDCs by following these steps:

1. In the **Cloud Resources** window, click **Organization VDCs**.

You can also click **Organizations** (in the **Cloud Resources** window), then drill down to the organization that contains the organization vDC.

2. Select one or more organization vDCs to back up and click one of the following backup tasks:

- To start an immediate organization vDC backup to server storage, click **Back Up Now**.
- To create a scheduled organization vDC backup to server storage, click **Create Backup Schedule**.

3. In the wizard, complete the following steps:

- a) Select the backup type:

Incremental Forever - Incremental (Default)

Backs up the blocks that changed since the previous backup (full or incremental). If a full backup does not exist for a vApp in this organization vDC, a full backup is automatically started. As a result, you do not have to verify that a full backup exists. After the initial full backup is

taken, an ongoing (forever) sequence of incremental backups occurs. This strategy requires no additional backup tasks to be defined.

Incremental Forever - Full

Creates an image of all vApps in this organization vDC. After the full backup is taken, there is no requirement to schedule more full backups.

- b) Select the data mover node that runs the backup operation. Select a data mover node that is not currently used in a backup or restore operation.
- c) If you clicked **Create Backup Schedule**, enter a name to identify this task. The task name must not contain any spaces. You can add a description for the task. If you clicked **Back Up Now**, you can change the default backup name (BackUpNow) and also add a description.
- d) Click **Next**.
 - If you clicked **Back Up Now**, click **Next** and proceed to Step 5.
 - If you clicked **Create Backup Schedule**, click **Next**, and proceed to Step 4.
4. In the **Schedule** page of the **Create Schedule** wizard, complete the following steps:
 - a) Specify the date and time to run the first backup.
 - b) Specify the interval that you want the backup to run.
5. Review the **Summary** page. If the information reflects your backup objective, click **Finish** to start the backup task or to create the schedule. Otherwise, click **Back** to make revisions.

Backing up data by disk usage

Specify the virtual machine disks that you want to include or exclude for backup services by setting include and exclude options.

Before you begin

Review the data mover `domain.vmfull`, `include.vmdisk`, and `exclude.vmdisk` options in [Client options reference](#).

About this task

To include virtual machines in your full virtual machine image backup operations, use the `domain.vmfull` option.

To include a virtual machine disk in a Backup VM operation, use the `include.vmdisk` option.

To exclude a virtual machine's disk from a Backup VM operation, use the `exclude.vmdisk` option.

Use these options for virtual disks that do not require backup. For example, use the options for those virtual disks that contain data that does not need to be restored, or the data is preserved by another backup mechanism.

Restriction: A virtual disk excluded from the backup operation is considered as deleted from the VM for that backup. If the VM is restored from that backup, the excluded virtual disk is not restored. Only the disk definition is restored.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

- Optional: View the disk name and label of the virtual disk by issuing the preview option.
For example:

```
dsmc backup vm VM1 -preview
```

- Exclude a virtual disk, set the `exclude.vmdisk` option in the IBM Storage Protect data mover `dsm.opt` options file.
For example:

```
EXCLUDE.VMDISK VM1 "Hard Disk 3"
```

- Issue the backup command:

```
dsmc backup vm VM1
```

Scenario: Including four disks for backup processing

Use the `include.vmdisk` and `domain.vmfull` options to include four virtual machine disks for backup services.

About this task

In the following examples, virtual machine `vm5_fin_com` contains four disks:

```
Hard Disk 1  
Hard Disk 2  
Hard Disk 3  
Hard Disk 4
```

Procedure

- Start a data mover command-line session:

- Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

- Use the `include.vmdisk` statement to back up disks Hard Disk 1 and Hard Disk 2.
For example:

```
INCLUDE.VMDISK vm5_fin_com "Hard Disk 1"  
INCLUDE.VMDISK vm5_fin_com "Hard Disk 2"
```

- Issue the backup command:

```
dsmc backup vm vm5_fin_com
```

Because an include disk statement is specified, this statement implies that only disks specifically included are backed up. As a result, Hard Disk 3 and Hard Disk 4 are not backed up.

- Use the `domain.vmfull` statement to back up disks Hard Disk 1 and Hard Disk 2:
For example:

```
DOMAIN.VMFULL "vm5_fin_com:vmrk=Hard Disk 1:vmrk=Hard Disk 2"
```

- Issue the backup command:

```
dsmc backup vm vm5_fin_com
```

Hard Disk 3 and Hard Disk 4 are not backed up.

You can include or exclude one or more disks with a `domain.vmfull` statement. You can specify include and exclude on the same statement. For example, the following statement is valid:

```
domain.vmfull "vm5_fin_com:vmdk=Hard Disk 1:-vmdk=Hard Disk 2:vmdk=Hard  
Disk 3:vmdk=Hard Disk 4"
```

If an include statement is present, it causes all other disks in the virtual machine to be excluded from a backup operation, unless the other disks are also specified with an include statement. For example, the following statement excludes all disks except for Hard Disk 1:

```
domain.vmfull "vm5_fin_com:vmdk=Hard Disk 1"
```

Scenario: Excluding four disks for backup processing

Use the `exclude.vmdisk` and `domain.vmfull` options to exclude four virtual machine disks for backup services.

About this task

In the following examples, virtual machine `vm5_fin_com` contains four disks:

```
Hard Disk 1  
Hard Disk 2  
Hard Disk 3  
Hard Disk 4
```

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Use the `exclude.vmdisk` statement to back up disks Hard Disk 1 and Hard Disk 2.

For example:

```
EXCLUDE.VMDISK vm5_fin_com "Hard Disk 3"  
EXCLUDE.VMDISK vm5_fin_com "Hard Disk 4"
```

3. Issue the backup command:

```
dsmc backup vm vm5_fin_com
```

Because an exclude disk statement is specified, this statement implies that only disks specifically excluded are not backed up. As a result, Hard Disk 3 and Hard Disk 4 are not backed up.

4. Use the `domain.vmfull` statement to back up disks Hard Disk 3 and Hard Disk 4:

For example:

```
DOMAIN.VMFULL "vm5_fin_com:vmdk=Hard Disk 3:vmdk=Hard Disk 4"
```

5. Issue the backup command:

```
dsmc backup vm vm5_fin_com
```

Hard Disk 3 and Hard Disk 4 are not backed up.

Scenario: Separating disks for backup and restore processing

To protect your data, coordinate the backup and restore capability of Data Protection for VMware and an IBM Data Protection agent installed in a guest virtual machine.

Before you begin

IBM Storage Protect provides applications that protect specific database and mail server data. The data protection application servers typically run in a virtual machine. To use Data Protection for VMware effectively with the IBM Storage Protect data protection applications, you must coordinate the backup and restore processing for each application. One way to coordinate backup and restore processing for each application is to separate processing by disk usage.

About this task

In this scenario, virtual machine VM2-08R2EX10-1 has IBM Storage Protect for Mail: Data Protection for Microsoft Exchange Server installed on Disk1 and uses this disk (.vmdk) configuration:

- Disk1: Operating system files
- Disk2: Microsoft Exchange Server database files
- Disk3: Microsoft Exchange Server log files
- Disk4: Contains files other than Microsoft Exchange Server files

Procedure

1. Use Data Protection for Microsoft Exchange Server to back up Disk2 and Disk3:

- a) Start a Data Protection for Microsoft Exchange Server command-line session by opening a command prompt and changing to the installation directory:
For example:

```
cd "C:\Program Files\IBM\StorageProtect\tdpexc"
```

- b) Issue the following backup command:

```
tdpexcc backup * full /backupmethod=vss /backupdestination=tsm
```

2. Use the data mover to back up Disk1 and Disk4:

- a) Start a data mover command-line session by opening a command prompt and changing to the data mover installation directory:
For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- b) Issue the following backup command:

```
dsmc backup vm "VM2-08R2EX10-1_orig:vmdk=Hard Disk 1:vmdk=Hard Disk 4"
```

3. Use the data mover to restore virtual machine VM2-08R2EX10-1 to a new virtual machine:
Issue the following restore command:

```
dsmc restore vm VM2-08R2EX10-1_orig -vmname=user_story_Exch  
-datacenter=VC4-VE-2_DATACENTER1 -host=ESX41-VE-2.QA1.COM  
-DATASTORE=ESX41-VE-3
```

The following output is displayed.

```
Restore processing finished.
Session established with server VM-03R2-TSM62-5: Windows
  Server Version 7, Release 1, Level 2.0
  Server date/time: 12/16/2014 12:32:54  Last access: 12/16/2014 11:13:13
```

```
Total number of objects restored:          2
Total number of objects failed:            0
Total number of bytes transferred:        42.00 GB
Data transfer time:                       4,708.17 sec
Network data transfer rate:               9,353.97 KB/sec
Aggregate data transfer rate:             9,210.25 KB/sec
Elapsed processing time:                   01:19:41
```

4. Start the VMware vSphere Client and go to **Summary > Edit Settings** to verify that the restored virtual machine (user_story_Exch) contains the same configuration as the original virtual machine (VM2-08R2EX10-1_orig). In this example, the configuration of the restored virtual machine shows four disks like the original virtual machine. However, only the data for Disk1 and Disk4 are restored. Disk2 and Disk3 must first be formatted for use. Then use Data Protection for Microsoft Exchange Server to restore Disk2 and Disk3.
5. Start the restored virtual machine:
 - a) Go to **Server Manager > Disk Management**.
 - b) Select Disk2 and Disk3 to complete disk initialization requirements.
 - c) Select partition style MBR and click **OK**.
 - d) Both disks are formatted with the same drive letters as the original virtual machine.
 - e) Use Data Protection for Microsoft Exchange Server to restore the Exchange data files to Disk2 and Disk3.

Backing up virtual machines by domain level

Narrow the focus of an operation to a subset of the virtual machines that are running on the system by setting the `domain.vmfull` option.

Before you begin

The `domain.vmfull` option backs up the virtual machines that are running on the system that is identified by the `vmchost` option.

Review the data mover `domain.vmfull` option. For more information, see [Domain.vmfull](#).

About this task

Complete these steps on the data mover system:

Procedure

1. Start a data mover command-line session:
 - **Windows** Open a command prompt and change to the data mover installation directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
 - `t_ve_bup_multiplevmsdm`
2. Open the data mover options file (`dsm.opt`) with your preferred text editor.
3. Enter the option name and one or more blank spaces, followed by the option value.
For example:

```
domain.vmfull vmhostcluster=Dev0105
```

Scenario: Backing up virtual machines by cluster server

Use the `domain.vmfull vmhostcluster` option to back up virtual machines for a specific cluster server.

About this task

The VMware environment consists of more than 3,000 VMs. Host clusters are used to manage the hardware resources. Although most of the clusters each contain 3 - 4 ESXi hosts, some clusters contain as many as 10 ESXi hosts. However, some ESXi hosts are running 1 - 3 VMs for larger, critical VMs. To manage the load, additional VMware hosts are frequently added or removed as they are needed. Each ESXi host in the cluster manages 10 - 30 VMs. Because the total number of VMs in each cluster ranges from 100 - 200, each host cluster is backed up to a dedicated vStorage backup server. Each server uses a dedicated data mover node to back up data.

Note: The Data Protection for VMware vSphere GUI does not support more than one VMware cluster with the same name under a single vCenter.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Include all virtual machines in cluster server TivDev01 in full VM backup operations.

- a) Set the `domain.vmfull` option with the **vmhostcluster** parameter in the client options file (`dsm.opt`).

For example:

```
domain.vmfull vmhostcluster=TivDev01
```

- b) Issue the backup command.

For example:

```
dsmc backup vm -vmbackuptype=fullvm
```

3. Repeat Step 2 for each cluster server.

Scenario: Backing up virtual machines by VMware datastore

Use the `domain.vmfull vmdatastore` option to back up virtual machines for a specific VMware datastore.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Include all virtual machines in VMware datastore `datastore_03` in full VM backup operations.

- a) Set the `domain.vmfull` option with the **vmdatastore** parameter in the client options file (`dsm.opt`).
For example:

```
domain.vmfull vmhostcluster=datastore_03
```

- b) Issue the backup command.
For example:

```
dsmc backup vm -vmbackuptype=fullvm
```

3. Repeat Step 2 for each datastore.

Scenario: Backing up virtual machines by name pattern

Use the `domain.vmfull vm` option to back up virtual machines by a specific name pattern.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Include all virtual machines that contain `MailDept` at the beginning of their name in full VM backup operations.

- a) Set the `domain.vmfull` option with the **vm** parameter in the client options file (`dsm.opt`).
For example:

```
domain.vmfull vm=MailDept*
```

- b) Issue the backup command.
For example:

```
dsmc backup vm -vmbackuptype=fullvm
```

3. Repeat Step 2 for each name pattern.

Backing up multiple virtual machines in parallel (optimized backup)

With parallel backup processing, you can use a single data mover node to back up multiple virtual machines (VMs) at the same time to optimize your backup performance.

Before you begin

To back up VMware VMs, the following options are provided so you can optimize the backups without adversely affecting the ESXi servers that are hosting the VMs. The options are described in detail in the [options reference](#).

vmmaxparallel

The `vmmaxparallel` option is used to control the maximum number of VMs that can be backed up at any one time. The optimal value for `vmmaxparallel` is not obvious; it depends on the processing power of the vStorage server that the data mover node runs on, and the performance of I/O between the data mover and the IBM Storage Protect server. For example, if you are moving data to the server over a busy LAN, you might need to limit the number of VMs in each parallel backup operation. Similarly, if the vStorage server processing capabilities are limited, for any reason, this is also a reason to restrict the value for `vmmaxparallel`.

vmmaxbackupsessions

The `vmmaxbackupsessions` option is used to control the maximum number of data movement sessions that can be included in the backup operation at any one time. Although this option sets the maximum number of sessions that are allowed, the `datamover` determines the actual number of sessions that are required based on the incoming workload and will use that number.

The value of the `vmmaxbackupsessions` option must be equal to or greater than the value of the `vmmaxparallel` option. If the value is less than the value of the `vmmaxparallel` option, a message is returned and the value is changed to the same value as `vmmaxparallel` option to ensure that there are as many sessions as there are VMs.

You might have to experiment with this setting to find the optimum value. Each dispatched VM is guaranteed one session and then extra sessions are applied to the dispatched VMs. The number of sessions will not exceed the value that is specified by the `vmmaxbackupsessions` option.

Other considerations for using this option include:

- If you are using the HotAdd data transport method, you will get better scale per session than if you are using network block device (NBD) data transports. This difference allows a higher value for the `vmmaxbackupsessions` option relative to a low value for the `vmmaxparallel` option. If you are using NBD transport, the difference between the `vmmaxbackupsessions` and `vmmaxparallel` options should be less because of scaling issues caused by having multiple NBD sessions per VM.
- There is no performance benefit for setting the `vmmaxbackupsessions` option if your storage system performance is slower than the available network speed between the data mover and the server.

vmmlimitperhost

The `vmmlimitperhost` option is used to control how many VMs and virtual disks can be backed up from an ESXi host at the same time.

You might have to experiment with this setting to find the optimum value. On ESXi servers that are heavily used, you might need to restrict the value for `vmmlimitperhost` so you do not adversely affect the vSphere server performance. On servers that are not as heavily used, you can include more VMs.

If you are using the NBD data transport method, you might also exceed the network file copy (NFC) protocol limit on the host if the value for `vmmlimitperhost` is too high. In this situation, a memory allocation error is returned as shown in the following example:

```
ANS9365E  VMware vStorage API error for virtual machine 'VM1'.
IBM Storage Protect function name : VixDiskLib_Read
IBM Storage Protect file          : ..\..\common\vm\vmvddksdk.cpp (3062)
API return code      : 2
API error message    : Memory allocation failed. Out of memory.
```

vmmlimitperdatastore

The `vmmlimitperdatastore` is used to control how many VMs and virtual disks can be backed up from a datastore at the same time. In a multiple datastore VMware environment, you can use this option to reduce the burden that is placed on any one datastore during a parallel backup operation.

Procedure

Complete these steps on the data mover system:

1. Start a command-line session:

- **Windows** Open a command prompt and change to the directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
 - **Linux** Open a terminal window and change to the directory: `cd /opt/tivoli/tsm/client/ba/bin`.
2. Open the `dsm.opt` file with your preferred text editor.
 3. Enter the option name and one or more blank spaces, followed by the option value.
For example:

```
vmmaxparallel 5  
vmmaxbackupsessions 10  
vmlimitperdatastore 5  
vmlimitperhost 5
```

4. Issue the **backup vm** command.
For example:

```
dsmc backup vm vm1 -vmbackuptype=fullvm
```

Using the examples provided, the backup operations for the VM `vm1` at the VM, virtual disk, or subdisk level can include 5 virtual machines and 10 sessions. Backup operations are limited to 5 VMs per datastore and 5 VMs per host.

Related information

[Backup VM](#)

Examples: Backing up multiple virtual machines in parallel

Parallel backup examples

In the following figures, the circled virtual machines are the virtual machines that are selected for backup processing, which is based on the option settings in `domain.vmfull`.

Example 1: Each VM is stored in a single datastore

Figure 5 on page 144 shows that each of the circled VMs has its data saved in a unique datastore. Assume that the parallel backup options are set to the following values:

- `vmmaxparallel 3`
- `vmmaxbackupsessions 3`
- `vmlimitperhost 1`
- `vmlimitperdatastore 1`




Figure 5. Virtual machines using unique datastores.

In Host A, only VMs 1 and 5 match the selection criteria on a `domain.vmfull` statement. In Host B, only VM 2 matches the selection criteria. In this configuration, each VM has a separate datastore, so the `vm_limit_per_datastore` setting is valid. But, since `vm_limit_per_host` is set to one, only one VM (vm1 or vm5) from Host A and one VM (vm2) from Host B are included when the Backup VM operation is run; only two VMs are included.

The `vm_max_backup_sessions` setting of 3 indicates that a backup operation for each of the two VMs will get a data movement session. Because there are three maximum backups sessions specified and only two VMs that are backed up, the backup operation for one of the VMs can get a second session. Sessions are obtained by the session pool manager.

Example 2: Same as example 1, but with a different setting for `vm_limit_per_host`

Figure 6 on page 145 shows that each of the circled VMs has its data saved in a unique datastore. In this configuration, the `vm_limit_per_host` is increased to two to illustrate how the option increase changes the Backup VM operation. Assume that the parallel backup options are now set to the following values:

- `vm_max_parallel 3`
- `vm_max_backup_sessions 3`
- `vm_limit_per_host 2` (an increase of 1)
- `vm_limit_per_datastore 1`




Figure 6. Virtual machines using unique datastores, with different option value for `vm_limit_per_host`.

The same VMs match the `domain.vmfull` criteria as they did in the previous example. However, with the increase in the `vm_limit_per_host` setting, now a total of three VMs are included in a Backup VM operation (vm1 and vm5 from Host A, and vm2 from Host B).

The `vm_max_backup_sessions` setting of 3 indicates that the backup operation for each of the three VMs will get a data movement session.

Example 3: Some VMs share datastores

Figure 7 on page 146 shows that the VMDK and configuration files for VM 5 in Host A is stored in two datastores. To include both vm1 and vm5 in Host A in the parallel backup operation, the value of `vm_limit_per_datastore` must be increased to at least two. If `vm_limit_per_datastore` is not increased to two, or higher, the backup of the second VM (vm5), in Host A, cannot be started until the first VM (vm1) backup is completed because the two VMs share data in data store1.




Figure 7. Virtual machines sharing a datastore.

Backing up virtual machines that host Active Directory controllers

About this task

The environment consists of a data center with five domain controllers (VDC1 - VDC5). The domain controllers are in two geographic locations. Each domain controller is on a VMware guest. One physical domain controller is included. The domain controllers are in two geographic locations and replicate by using an Active Directory replication process.

Procedure

1. Start a data mover command-line session:

Windows Open a command prompt and change to the data mover installation directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.

2. Back up the virtual machine guests that host VDC1 and VDC3.
In these examples, virtual machine guest VM1 hosts domain controller VDC1, and virtual machine guest VM3 hosts domain controller VDC3:

```
dsmc backup vm VM1
```

```
dsmc backup vm VM3
```

3. Recover the virtual machine guest that hosts VDC1.
In this example, virtual machine guest VM1 hosts domain controller VDC1:

```
dsmc restore vm VM1
```

The domain controller is restored to the version taken at the time of the backup. After the server restarts, the restored Active Directory domain controller (on VDC1) replicates data by using other domain controllers in the network.

4. Restart the restored virtual machine guest.
5. Verify that replication processing completed successfully.

Related tasks

“[Verifying that the Active Directory Domain Controller replicated successfully](#)” on page 158

When a VM guest that contains an Active Directory Domain Controller (AD DC) is restored with Data Protection for VMware, the DC (on that VM) is restored from a backup copy of the AD database.

Specifying a management class to associate objects

Specify how to manage virtual machine and vApp backups operations on the IBM Storage Protect server.

Before you begin

Review the data mover `vmmc`, `vmctlmc`, and `vappmc` options in [Client options reference](#).

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Open the data mover options file (`dsm.opt`) with your preferred text editor.
3. Enter the option name and one or more blank spaces, followed by the option value.
For example:

```
vmmc myManagementClass
```

```
vmctlmc diskonlymc
```

```
vappmc MCPRDVPAPS
```

Scenario: Specifying a management class for VMware backups in a vSphere environment

Use the `vmmc` option to store the VMware backups with a management class other than the default management class.

About this task

Assign a virtual machine backup to a non-default management class.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Create a full VM backup of the virtual machine.

Assign the backup to a non-default management class.

For example, to assign the backup of virtual machine `myVirtualMachine` to management class `myManagementClass`, specify the `vmmc` option in the command. For example:

```
dsmc backup vm "myVirtualMachine" -vmmc=myManagementClass
```

For information about how management class rebinding applies to VMware backup data on the IBM Storage Protect server, see [technote 1665032](#).

Scenario: Specifying a management class for VMware control files in a vSphere environment

Use the `vmctlmc` option to assign the VMware control files to a management class other than the default management class.

Before you begin

VMware control files are assigned to the default management class. Use the `vmmc` option to assign VMware data and VMware control files to a non-default management class. The `vmctlmc` option overrides the default management class and the `vmmc` option for VMware control files.

Under certain conditions, it might be necessary to assign VMware control files to a different management class than the VMware data files.

Use the `vmctlmc` option if VMware data files are backed up to tape. Back up the VMware control files to a disk-based storage pool that is not migrated to tape. The storage pool can contain random access volumes and sequential file volumes; the storage pool can be a deduplicated pool. Use the `vmctlmc` option to specify a management class that stores control files in such a storage pool.

Restriction: The management class that is specified by the `vmctlmc` option determines only the destination storage pool for VMware control files. Retention of VMware control files is determined by the `vmmc` option, if specified, or by the default management class. The retention for the VMware control files always matches the retention of the VMware data files.

About this task

Assign a virtual machine backup to a non-default management class.

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- **Linux** Open a terminal window and change to the data mover installation directory. For example:

```
cd /opt/tivoli/tsm/client/ba/bin
```

2. Create a full VM backup of the virtual machine.

Assign the backup to a non-default management class.

For example, to assign the backup of virtual machine `myVirtualMachine` to management class `myManagementClass`, specify the `vmmc` option in the command:

```
dsmc backup vm "myVirtualMachine" -vmmc=myManagementClass
```

Specifying objects to include in backup and restore operations

Specify the VMs or vApps that you want to include for backup services by setting include options.

Before you begin

Review the data mover `include.vm` and `include.vapp` options. For more information, see [Virtual machine include options](#).

About this task

Complete these steps on the data mover system:

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
- **Linux** Open a terminal window and change to the data mover installation directory: `cd /opt/tivoli/tsm/client/ba/bin`.

2. Open the data mover options file (`dsm.opt`) with your preferred text editor.

3. Enter the option name and one or more blank spaces, followed by the option value.
For example:

```
include.vm vmtest*
```

Scenario: Specifying objects to include for backup and restore operations in a vSphere environment

Use include options to specify the objects that you want to include in backup and restore operations.

Before you begin

To define when objects are included in a backup, how long they are kept on the server, and how many versions of the object the server keeps, use a management class. Set the management class for an object by using the data mover **vmmc** option. Place this option in the client options file `dsm.opt` or in the client system options file `dsm.sys`.

You can also change how files are processed, for example to use a different management class, by using the data mover **include.vm** option.

About this task

This scenario assumes the following active management classes on the IBM Storage Protect server:

- MCFORTESTVMS
- MCFORPRODVMS
- MCUNIQUEVM

Procedure

1. Start a data mover command-line session:

- **Windows** Open a command prompt and change to the data mover installation directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
- **Linux** Open a terminal window and change to the data mover installation directory: `cd /opt/tivoli/tsm/client/ba/bin`.

2. Open the data mover options file (`dsm.opt`) with your preferred text editor.

3. Associate all virtual machine backups, with names that begin with VMTEST, to management class MCFORTESTVMS:

```
include.vm vmtest* MCFORTESTVMS
```

4. Associate virtual machine backup WINDOWS VM1 [PRODUCTION] to management class MCFORPRODVMS:

```
include.vm "WINDOWS VM1 ?PRODUCTION?" MCFORPRODVMS
```

The following values are used:

- The virtual machine is enclosed in quotation marks because it contains space characters.
- The question mark (?) wildcard is used to match the special characters in the virtual machine name.

5. Associate virtual machine backup VM1 to management class MCUNIQUEVM:

```
include.vm VM1 MCUNIQUEVM
```

Chapter 7. Restoring VMware data

Data Protection for VMware restore scenarios are provided in this section.

Mounting a virtual machine disk and exporting the volumes

You can restore one (or more) files from a virtual machine that was backed up to IBM Storage Protect server storage.

Before you begin

File restore from tape media is not supported. File restore from disk storage is the preferred method.

Consider moving target virtual machine backup data from tape media to disk storage before you attempt a file restore operation. You can move data with the server **MOVE NODEDATA** command. You can also run traditional full VM backups regularly.

Linux The mount operation fails if the VM contained one or more physical volumes that did not have a volume group assigned at the time of back up.

Procedure

To mount a backed up virtual machine disk and export the mounted volume for a file restore operation, complete the following steps:

1. Configure the mount proxy nodes:

- a) Go to the **Configuration** window in the Data Protection for VMware vSphere GUI.
- b) Click **Edit IBM Storage Protect Configuration** in the **Tasks** list. The configuration notebook might take a few moments to load.
- c) Go to the **Mount Proxy Node Pairs** page and select a VMware datacenter in the table.
- d) Click **Add Mount Proxy Node Pair**.
- e) Click **New Settings** in the table.

For Linux operating systems, the Linux mount proxy node must be configured manually. Use the sample `dsm.sys` file content that is shown in the **Mount Proxy Settings** dialog when you configure the Linux mount proxy node.

For Windows operating systems, only one client acceptor is created. If you want to add a second Windows mount proxy node, you must manually configure the client acceptor on a remote system.

- f) Specify the storage device type from which the snapshot is mounted by setting the client `vmstoragetype` option in the client options file on the Windows mount proxy node.

- i) Open a command prompt and change to the data mover installation directory. For example:

```
cd "C:\Program Files\IBM\StorageProtect\baclient"
```

- ii) Open the data mover options file (`dsm.opt`) with your preferred text editor.

- iii) Set the `vmstoragetype` option with one of the following device types:

DISK

The snapshots to be mounted are on Disk or File storage pools. This value is the default.

VTL

The snapshots to be mounted are on VTL storage pools.

TAPE

The snapshots to be mounted are on Tape storage pools.

2. Validate that the mount proxy nodes are online and that the iSCSI service is running:

- a) Go to the **Configuration** window in the Data Protection for VMware vSphere GUI.
 - b) Select a VMware datacenter in the navigation tree.
 - c) Select the mount proxy node (created in Step 1) in the table and click **Validate Selected Node**. The validation results might take a few moments to complete.
A successful validation shows the Status: Running message for each mount proxy node in the **Status Details** pane. An unsuccessful validation shows the Status: error message for each mount proxy node that encountered an error.
Remember: By default, the iSCSI service is started manually. If the system where this service is running restarts, the iSCSI service must be restarted manually.
If you receive an error message, investigate possible causes that are based on error messages that are shown in the **Task Details** pane of the **Recent Tasks** report.
3. Go to the **Restore** window in the Data Protection for VMware vSphere GUI and click **Restore Points** in the header.
 4. Expand the list of VMware datacenters and select a virtual machine from the navigation pane. All active and inactive backup versions for the selected virtual machine are identified as restore points in the **Restore Points** table. Virtual machine template backups are identified in the **Template** column in the **Restore Points** table.
Depending on the number of managed vCenters, the list might take a few moments to load.
 5. Select one or more restore points for one virtual machine in the table and click **Mount** to open the **Mount** wizard.
 6. On the **Mount Options** page, complete the following steps:
 - a) Select the mount proxy node.
 - b) Optional: To select the operating system where the backed up virtual machine disks are to be mounted, click the **Guest Operating System** list and make your selection.
Tip: When the operating system where the disks are to be mounted is different from the operating system of the mount proxy node, the mount point path updates automatically.
 - c) Enter the absolute path of the mount point. The following characters are supported: a-z, A-Z, 0-9, colon (:), forward slash (/), backward slash (\), and underscore (_). The maximum length is 200 characters.
 - **Windows** The absolute path to a disk is *mount root\tag\vmname\snapshot date and time\file system number*. For example:

```
C:\tsmmount\ticket_9471\tangowin2k12test\2014-07-01-10_35_50\Volume1\
```


The default value is *C:\tsmvemount\vmname*.
Restriction: The maximum length of the path and file name that is accessed in a mounted volume on Windows cannot exceed 6255 characters. This maximum length includes the total characters in the path, volume name, mount point, virtual machine name, tag description, and snapshot date.
 - **Linux** For Linux operating systems, the absolute path to a disk is *mount root/tag/vmname/snapshot date and time/file system number*. For example:

```
/tsmmount/ticket_9471/vm1/2014-07-01-10_35_23/Volume1
```


The default value is */mnt/vmname*.
Restriction: The maximum length of the path and file name that is accessed in a mounted volume on Linux cannot exceed 4096 characters. This maximum length includes the total characters in the path, volume name, mount point, virtual machine name, tag description, and snapshot date.
 - d) Enter a description of this mount operation in the **Description Tag** field. This description becomes part of the mount path so that the administrator can easily identify the operation. The following characters are supported: a-z, A-Z, 0-9, and underscore (_). The maximum length is 20 characters.

- e) Optional: If you want the mounted virtual machine disks on a network share, select **Create Network share** and enter the appropriate credentials.
 - For Windows operating systems, enter the user name that is allowed to access Windows Share.
Tip: For security reasons, ensure that files are shared only by users and groups who are authorized to access those files.
 - For Linux operating systems, enter the IP address or name of the system that mounts the exported file system.
- f) Click **Next**.
- 7. In the **Summary** page, review the settings and click **Finish** to start the mount operation. To change your mount settings, click **Back**.
After the operation starts, you can monitor its progress (**Recent Tasks**) in the **Reports** window.
If the mount operation does not complete successfully, investigate possible causes that are based on error messages that are shown in the **Task Details** pane of the **Recent Tasks** report.
- 8. Export (or share) the mounted volumes from the virtual machine:
 - a) Go to the **Restore** window in the Data Protection for VMware vSphere GUI.
 - b) Click **Mount Status** in the header.
 - c) Select the mount operation that contains the volumes you want to export.
 - d) Copy the content in the **Network Share** pane by using Ctrl + C and send to the user who accesses the mounted volumes to restore the files.
- 9. Log in to the system where the files will be restored and complete the following step:
 - For Windows operating systems, connect to the Common Internet File System (CIFS) where the files are mounted. Copy the files with a file manager application such as Windows Explorer.
 - For Linux operating systems, connect to the Network File System (NFS) where the files are mounted. Copy the files with a file manager application.

What to do next

After the files are restored by the user, dismount the volumes:

1. Go to the **Restore** window in the Data Protection for VMware vSphere GUI.
2. Click **Mount Status** in the header.
3. Select the mount operation that contains the volumes you want to export and click **Dismount**. Your mount operation is identified by Type=HelpDesk in the **Mount Status** table.

vSphere environment restore scenario

This scenario demonstrates how to restore VMs with the **vmcli -f restore** command.

The VMs can also be restored with the following user interfaces:

Data Protection for VMware vSphere GUI

Information about how to complete restore tasks with the Data Protection for VMware vSphere GUI is provided in the online help that is installed with the GUI. Click **Learn More** in any of the GUI windows to open the online help for task assistance.

IBM Storage Protect backup-archive client GUI

Information about how to complete restore tasks with the backup-archive client GUI is provided in the online help that is installed with the GUI. Click **Help** in any of the GUI windows to open the online help for task assistance.

IBM Storage Protect backup-archive client command-line interface

Information about how to complete restore tasks with the **dsmc Restore VM** command is provided in the command-line help that is installed with the product (**dsmc help restore vm**). Information is also available at [Restore VM](#).

This scenario completes an instant restore of vmName6 to a new VM, data center, ESX host, and data store. During the restore process, the disk is created with thin provisioning.

The following **vmcli -f restore** command is issued:

```
vmcli -f restore -vmrestoretype instantrestore -I vmlistfile
```

The **vmplistfile** contains this statement:

```
backupid:678912345 vmname:vmName6::vmname:vm6newName  
newdatacentername:DataCenter2 newesxhostname:esxHost1Name  
newdatastoreurl:datastore2 vmtempdatastore:datastore2temp  
vmdiskprovision:thin
```

Windows Full VM instant restore scenarios

A backed up VM is restored and available for immediate use.

Instant access and instant restore capability is supported only for VMware VMs that are hosted on VMware ESXi 5.1 servers, or later versions.

This command completes an instant restore of the VM with the name Cologne.

```
dsmc restore vm Cologne -vmrest=INSTANTRestore -vmtempdatastore=Verify_datastore
```

This command completes a regular restore (without starting the VM) when the VM named San_Jose is being restored.

```
dsmc restore vm San_Jose
```

Alternatively, you can also use the following command:

```
dsmc restore vm San_Jose -vmrest=noni
```

In this command, the **-vmtempdatastore** option specifies a temporary data store on the ESX host. The data for the new VM is stored in this temporary data store:

```
dsmc restore vm Haifa -VMRESToretype=INSTANTRestore -vmname=Haifa_verify  
-VMTEMPDatastore=Verify_Datastore
```

The temporary data store is used by Storage vMotion to store the configuration of the restored virtual machine during the instant restore process. The name that you specify must be unique. It cannot match the name of any of the original data stores that were used by the virtual machine when it was backed up, and it cannot be the same as the name specified on the optional **-datastore** option. If the **-datastore** option is omitted, the virtual machine files are restored to the data stores that they used when the virtual machine was backed up.

This command completes an instant restore of the VM with the name Oslo with the **-pick** option to pick a specific backup version.

```
dsmc restore vm Oslo -vmrest=INSTANTRestore -pick -vmtempdatastore=datastore_temp
```

This command queries to find all active and failed instant restore sessions.

```
dsmc query vm * -VMRESToretype=INSTANTRestore
```

This command retrieves the metadata for all instant restore session from the IBM Storage Protect server and prints that information as a list.

This command completes a cleanup of the VM and all its components. These components include iSCSI mounts, devices, and temporary data that are identified by the VM name on the ESX host.

```
dsmc restore vm Oslo -VMRESToretype=VMCleanup -vmname=Oslo_Verify
```

This command queries information about the real state of the listed VMs based on information from the vSphere SDK on the ESX host.

```
dsmc query vm * -VMRESToretype=INSTANTRestore -Detail
```

Full VM instant restore cleanup and repair scenarios

When an instant restore operation fails after the VM is powered on, manual cleanup and repair tasks are required.

An instant restore operation that fails with storage vMotion running creates either of the following situations:

- The instant restore operation generates an error message.
- The instant restore operation suspends indefinitely and the VM is not responsive.

To determine the cause of the problem, perform a detailed query of the VM by using the following command:

```
dsmc q vm * -vmrestoretype=instantrestore -detail
```

In the output that is produced by this command, for each VM in the output, look for the line that contains **Action Needed**. Use the following *Action Needed* paragraphs to recover from failed instant restore operation, depending on the **Action Needed** status.

Instant access and instant restore capability is supported only for VMware VMs that are hosted on VMware ESXi 5.1 servers, or later versions.

Action Needed: Cleanup

In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the storage vMotion status is successful (**vMotion Status: Successful**) and that all VM disks are physical disks (**Disk Type: Physical**). This status confirms that the VM was restored and cleanup of orphaned components, such as iSCSI mounts, is needed.

This type of failure occurs as a result of either of the following situations:

- The instant restore failed and Storage vMotion is running. VMware vSphere continues the vMotion process.
- Storage vMotion finished successfully, but the automatic cleanup of the iSCSI mounts fails.

To clean up any orphaned components, run the **restore vm** command with the **-VMRESToretype=VMCleanup** parameter. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

Action Needed: Repair

In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the iSCSI device that is attached to the VM is dead (status is **Disk Path: Dead**).

This type of failure occurs as a result of one of the following three situations:

- The VM that is used as a data mover or the physical data mover machine failed.
- A network failure occurred between the data mover and the ESX host or the data mover and the IBM Storage Protect server.
- The IBM Storage Protect recovery agent Service failed.

The iSCSI device must be returned to an active state before any other instant operation is attempted.

To attempt to recover from a data mover failure, complete the following steps:

1. Investigate that cause of the failure and restart the data mover machine if it does not start automatically. This action starts an automatic recovery of the mounted iSCSI disks.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. Restart storage vMotion in the vSphere client and monitor its progress in the vSphere client status bar.
4. If storage vMotion processing completed successfully, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

To attempt recovery after a network failure, complete the following steps:

1. Repair the network issue so that communication between the data mover and the ESX host, and the data mover and the IBM Storage Protect server resumes.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. If the network failure did not cause storage vMotion to time out, no action is required.
4. If the network failure caused storage vMotion to time out, and the error message indicates that the source disk is not responding, restart storage vMotion in the vSphere client. When storage vMotion processing completes, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

To attempt recovery after a recovery agent service failure, complete the following steps:

1. Investigate that cause of the failure and restart the recovery agent service if it does not start automatically. This action starts an automatic recovery of the mounted iSCSI disks.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. If the recovery agent service failure did not cause storage vMotion to time out, no action is required.
4. If the recovery agent service failure caused storage vMotion to time out, and the error message indicates that the source disk as not responding, restart storage vMotion in the vSphere client. When storage vMotion processing completes, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

Full cleanup

If you are not able to recover from a failure and want to remove the VM and its components, run the **restore vm** with the **-vmrestoretype=VMFULLCleanup** parameter. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMFULLCleanup
```

A **VMFULLCleanup** operation forces removal of the VM and all of its components, regardless of the state of the virtual machine. Do not start a full clean up operation while vMotion is still migrating a virtual machine.

A new VM is built from the requested VM backup and is available for immediate use.

The process that creates a VM for verification is referred to as instant access. The verification itself is done by a specific application that the user must provide and operate. In this scenario, since the VM guest data remains in the IBM Storage Protect server repository, no additional storage is required on the primary data store.

Instant access and instant restore capability is supported only for VMware VMs that are hosted on VMware ESXi 5.1 servers, or later versions.

Start an instant access scenario

Start the instant access operation by entering the following IBM Storage Protect backup-archive client command and options:

```
dsmc restore vm Haifa -VMRESToretype=instanta -vmname=Haifa_verify
```

This command verifies the VM backup named Haifa without having to restore the VM. Since the original VM exists, the `-vmname` option assigns the new VM name Haifa_verify.

In this command, the `-vmautostart` option specifies that the VM created for verification is powered on automatically:

```
dsmc restore vm Haifa -VMRESToretype=instanta -vmname=Haifa_verify  
-VMAUTOSTARTvm=YES
```

By default, the VM created for verification is not powered on automatically. This default value allows the VM to be reconfigured before startup (if needed).

Specify the `-inactive` and `-pick` options to select the VM backup to validate from a list of all backed up VMs. Or, specify the `-pitdate` and `-pittime` options to select a VM backup by its backup date and time.

All location options (such as `-vmname`, `-datacenter`, `-host`, and `-datastore`) are supported by the `-vmrestoretype=instantaccess` and `-vmrestoretype=instantrestore` options.

This command returns a list of VMs that are running in instant access mode:

```
dsmc query vm * -VMRESToretype=instanta
```

This command starts the cleanup process for a VM backup:

```
dsmc restore vm Haifa -VMRESToretype=VMCleanup -vmname=Haifa_Verify
```

This command completes the instant access session. A cleanup process includes these actions:

- The temporary VM created for verification is deleted on the ESX host.
- The iSCSI mounts are unmounted on the vStorage Backup Server.
- The iSCSI device list is cleaned on the ESX host.
- The temporary data that is generated during verification is deleted by removing the VMware snapshot.

You cannot use the `-VMRESToretype=VMCleanup` option or the `-VMRESToretype=VMFULLLCleanup` option to clean up a virtual machine while Storage vMotion is migrating it to a host.

Verifying that the Active Directory Domain Controller replicated successfully

When a VM guest that contains an Active Directory Domain Controller (AD DC) is restored with Data Protection for VMware, the DC (on that VM) is restored from a backup copy of the AD database.

Before you begin

The original VM must be powered off before the restored VM is started. In addition, the restored VM must be manually rebooted for replication to occur.

About this task

The following tasks occur upon a successful Data Protection for VMware restore and subsequent reboot of the VM guest that contains the AD DC:

Procedure

1. The DC is updated from a backup copy of the AD DC database. A new invocationID is assigned to the Directory Server. This update is indicated by event 1109 in the event log on the VM guest. To verify this update:
 - a) In the **Computer Management** window on the restored system, go to **System Tools > Event Viewer**.
 - b) When the AD DC restored successfully, the Information type event for the restored DC displays the following message:

```
ActiveDirectory 1109 Replication
```

The message in the Event Viewer also confirms a successful restore because of the changed invocationID attribute:

```
The invocationID attribute for this directory server has been changed.
The highest update sequence number at the time the backup was created is <time>
InvocationID attribute (old value):<Previous InvocationID value>
InvocationID attribute (new value):<New InvocationID value>
Update sequence number:<USN>
The InvocationID is changed when a directory server is restored from backup media
or is configured to host a writeable application directory partition.
```

2. The restored DC replicates itself non-authoritatively with its replication partners in the network. It is updated with the most current domain, schema, configuration, and application partitions:

Note: Data Protection for VMware does not support authoritative restore.

- a) Log in to the VM guest that was restored by using Data Protection for VMware as an Administrator.
- b) Open a Windows command prompt.
- c) Check the status of the last replication that involved the restored DC by issuing the `repadmin /showrepl` command¹. This command shows the replication partners for each directory partition on the DC and the status of the last replication.

If the replication schedule did not start, you can manually start the replication operation. Go to the **Active Directory Sites and Services**, select the replication partners, and right-click **Replicate Now**.

When the status is newer than the restore time, this status means that the replication was successful and completed automatically. The following output shows that replication was successful:

¹ Repadmin.exe is a Microsoft command-line tool that is installed with Microsoft Active Directory.


```

Repadmin: running command /showrepl against full DC localhost
Default-First-Site-Name\DC12012
DSA Options: IS_GC
Site Options: <none>
DSA Object GUID: 8393da24-f18b-453a-b197-b8dc6956d51f
DSA invocationID: 8393da24-f18b-453a-b197-b8dc6956d51f

==== INBOUND NEIGHBORS =====

CN=Configuration,DC=his,DC=local
  Default-First-Site-Name\DC22012 via RPC
    DSA Object GUID: 790c6f2d-61f1-4704-bdcf-6ef731bcb96e
    Last attempt @ 2013-01-25 14:33:10 was successful.

```

When the repadmin /showrepl command displays a successful replication, the AD DC replication is considered successful. No additional tasks are required.

- d) When the repadmin /showrepl command shows that replication was not successful, output similar to the following is shown:

```

Repadmin: running command /showrepl against full DC localhost
Default-First-Site-Name\DC12012
DSA Options: IS_GC
Site Options: <none>
DSA Object GUID: 8393da24-f18b-453a-b197-b8dc6956d51f
DSA invocationID: 8393da24-f18b-453a-b197-b8dc6956d51f

==== INBOUND NEIGHBORS =====

CN=Schema,CN=Configuration,DC=his,DC=local
  Default-First-Site-Name\DC22012 via RPC
    DSA Object GUID: 790c6f2d-61f1-4704-bdcf-6ef731bcb96e
    Last attempt @ 2013-01-25 14:30:32 failed, result 1908 <0x774>:
    Could not find the domain controller for this domain.
    1 consecutive failure(s).
    Last success @ 2012-12-14 15:01:36.

```

If a replication failure exists or persists, follow the instructions provided in the next section.

Recover from Replication Failures

Use the following methods to investigate the cause of a persistent replication failure:

1. Use the Microsoft Domain Controller Diagnostics tool (dcdiag.exe) to view information about all components, objects, and permissions that are required for successful replication. For example:
 - a. Open a Windows command prompt as an administrator.
 - b. Issue the dcdiag /test:replications command. Use the output information to resolve any issues. If the command fails, investigate the events that are at **Event Viewer > Directory Service > ActiveDirectory_DomainServices**.
2. Use the Microsoft Repadmin.exe command-line tool to view the retired invocationID on a DC. For example:
 - a. Open a Windows command prompt as an administrator.
 - b. Issue the repadmin /showsig [DC_LIST] command. This output shows that restore from the IBM Storage Protect server was successful because a retired invocationID exists:

```

C:\Users\Administrator>repadmin /showsig rodc
Default-First-Site-Name\RODC

Current DSA invocationID: ed8ea6b9-d347-4695-b886-b5128be280c4
2c995946-2389-4d98-bc78-3708ba906e01 retired on 2012-12-19 16:56:21
at USN 17703

```

When the output contains the statement No retired signatures, the AD was not restored from the server correctly. As a result, replication cannot be completed because the partner DCs mistake the new invocationID as evidence for a completed replication. For example:

```
C:\Users\Administrator>repadmin /showsig rodc
Default-First-Site-Name\RODC

Current DSA invocationID: ed8ea6b9-d347-4695-b886-b5128be280c4
No retired signatures
```

When the invocationID is retired, the replication can be started. However, this statement does not guarantee success of the replication.

Restoring a virtual disk using multiple sessions

To optimize performance for restore operations, multiple sessions can be used to restore a virtual disk.

Before you begin

To restore a virtual disk using multiple sessions, use the `vmmaxrestoresessions` option. This option specifies the maximum number of IBM Storage Protect server sessions that can be used in an optimized restore operation for a virtual disk.

Multiple session restores can be used under the following conditions:

- The data to be restored is stored on several tape volumes or file device class volumes.
- Sufficient mount points are available.

Note: This option is not valid for Hyper-V virtual machines backups.

About this task

Complete these steps on the data mover system:

Procedure

1. Start a command-line session:

- **Windows** Open a command prompt and change to the directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
- **Linux** Open a terminal window and change to the directory: `cd /opt/tivoli/tsm/client/ba/bin`.

2. Open the `dsm.opt` file with your preferred text editor.

3. Enter the `vmmaxrestoresessions` option and one or more blank spaces, followed by the option value.

For example:

```
vmmaxrestoresessions 3
```

4. Issue the **restore vm** command.

For example:

```
dsmc restore vm vm1
```

Using the examples provided, the restore operations for virtual disks in the VM `vm1` can use a maximum of 3 sessions.

Note: The maximum number of sessions that can be allocated to an individual disk being restored is based on the transport the disk was opened with. The following are the limits of a restore sessions number:

- For NBD and NBDSSL: 2 sessions
- For HOTADD and SAN: 4 sessions

Related information

[Restore VM](#)

Restoring virtual disks using a CSV file

To specify different options for parallel restore operations, you can use a CSV file.

Before you begin

Command-line restore provides a wide range of options that are applied simultaneously to a number of virtual machines. This works well for a large number of machines using the same settings, but it does not provide flexibility in supporting parallel restore operations with different settings.

As an alternative to the command-line syntax, you can use a comma separated values (CSV) file to define the list of virtual machines to restore, using a subset of the available options (-vmname, -datastore, -datacenter, -host, -pitdate, and -pittime). Using a CSV file provides more granularity than pure command-line syntax and can deal with more complex situations.

CSV files can be output from standard spreadsheet applications, providing ease-of-use and avoiding complex command line syntax. You can add one or more columns to store notes and observations: only the column headings defined below will be recognized and acted on. In some versions of vSphere Web client you can generate a CSV file from any list view. In the Inventory List, you can use the Export List functionality to create a CSV report listing all VMs in a host cluster. This output can provide a starting point for the CSV file you use to drive bulk restore operations.

Each row in the CSV after the header row represents a virtual machine to be restored. Unlike pure command-line parallel restore, each virtual machine may take a different set of options. You can specify a mix of active and inactive backups in the same restore operation. The following column heading names are recognized by the command.

Table 10. Column heading names

Heading	Description	Usage
Virtual Machine Name	The name of the virtual machine to be restored.	No wildcard characters are allowed. Case-sensitive. This column is mandatory.
New Virtual Machine Name	The name of the virtual machine that is restored.	This column uses the same syntax as the -vmname option. Optional. You can leave this column blank if you want to reuse the existing name.
New Datastore	The new datastore to which the virtual hard disks are restored.	This column uses the same syntax as the -datastore option. Optional. You can leave this column blank if you want to reuse the existing datastore.
New Datacenter	The new datacenter with which the virtual machine should be associated.	Uses the same syntax as the -datacenter option. Optional. You can leave this column blank if you want to reuse the existing datacenter.
New Host	The new host to which the virtual machine will be restored.	This column uses the same syntax as the -host option. Optional. You can leave this column blank if you want to reuse the existing host.

Table 10. Column heading names (continued)

Heading	Description	Usage
PITDATE	The point-in-time date from which the backup is specified.	This column uses the same syntax as the -pitdate option. Optional. You can leave this column blank to indicate the active backup should be restored. This column is required if PITTIME is specified in the CSV file. PITDATE dates should use the format set by the DATEFORMAT option. The default varies by locale in Windows. The default is DATEFORMAT 1 in Linux.
PITTIME	The point-in-time time of day from which the backup is specified.	This column uses the same syntax as the -pittime option. Optional. You can leave this column blank to indicate you want to use the active backup or if only the PITDATE is specified. PITTIME times should use the format set by TIMEFORMAT option. The default varies by locale in Windows. The default is TIMEFORMAT 1 in Linux.

Procedure

1. Start a command-line session:

- **Windows** Open a command prompt and change to the directory: `cd "C:\Program Files\IBM\StorageProtect\baclient"`.
- **Linux** Open a terminal window and change to the directory: `cd /opt/tivoli/tsm/client/ba/bin`.

2. Restoring from a CSV file uses the following syntax:

```
restore vm C:\restore-virtual-machines-list.csv -csv
```

Related information

[Restore VM](#)

[Vmmaxrestoreparallelvms](#)

Appendix A. Troubleshooting

Solutions to Data Protection for VMware vSphere GUI and Data Protection for VMware command-line interface issues are provided.

Locating log files

For information about Data Protection for VMware log files, see the following topics:

- [Log file activity](#)
- [File restore log activity options](#)
- [“Trace options for file restore” on page 169](#)

Data Protection for VMware vSphere GUI backup or restore operation fails

Complete these tasks to resolve a backup or restore failure:

1. Log on to the system where the data mover is installed.
2. Start a command-line session:
 - **Windows** Open the Windows **Start** menu and select **Programs > IBM Storage Protect > Backup-Archive Command Line**.
 - **Linux** Open a terminal window.
3. If not already there, go to the installation directory:

```
Windows cd C:\Program Files\IBM\StorageProtect\baclient
```

```
Linux cd /opt/tivoli/tsm/client/ba/bin
```

By default, error log files are in the installation directory.

4. View these data mover log files to see if an error was generated:
 - `dsmerror.log`: All client messages.
 - `dsmwebcl.log`: All web client messages.
 - `dsmj.log`: All client Oracle Java™ GUI 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 IBM IBM Documentation at [Messages, return codes, and error codes](#).

5. If neither of these files contain an error, run a backup-archive client backup and restore operation to see if it fails.
6. If the data mover operations complete successfully, run a Data Protection for VMware command-line interface [“Backup” on page 97](#) and [“Restore” on page 99](#) operation. Set the appropriate trace parameters (as described in [“Profile parameters” on page 115](#)) 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 data mover `dsmerror.log`:

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

Typically, this error results when the VMCLI node name is different from its administrator name. These two names must be the same.

Data mover nodes are not visible during a backup operation

Verify that the correct proxy node authority was granted on the IBM Storage Protect server. If the correct authority exists, then the data center mapping specified by the `VE_DATACENTER_NAME` profile parameter is incorrect. See [“Profile parameters” on page 115](#) for a complete description and correct syntax of the `VE_DATACENTER_NAME` parameter.

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 IBM Storage Protect 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.

Session timeout

The IBM Storage Protect 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 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.

Resolving a VM guest (with application protection) backup failure

In this situation, a backup (with application protection) of a guest machine is stopped by the user. When the data mover backup process (**dsmagent** or **dsmc**) ends in this manner, the cleanup of the application protection does not take place. As a result, the next backup (with application protection) of the same guest machine can be issued only after a 10-minute interval. This interval is the length of time necessary for the process to recognize that the guest machine is not backed up.

To manually clean up application protection without waiting 10 minutes for communication to clear, complete these steps:

1. Log on to the guest machine with the same user ID and password that was entered when you issued the backup operation.
2. Open a command prompt window and issue this command:

```
echo %TEMP%
```

3. Go to the %TEMP% directory, then change to the TSM directory. For example:

```
C:\Users\Administrator\AppData\Local\Temp\TSM
```

4. Delete the BackupHeartBeat.txt file.
5. Back up the guest machine.

Event log contains event ID 8194, VSS message

After a backup of a VM guest with application protection completes, the event log contains the event ID 8194, VSS error message. This cause of this message is an incorrect security setting in the Volume Shadow Copy Service (VSS) writer or requestor process.

To resolve this error, complete these steps:

1. Log on to the VM guest as an administrator and run the Microsoft dcomcnfg.exe utility in the **Start > Run** dialog:

```
dcomcnfg.exe
```

Click **OK**.

The dcomcnfg.exe utility is used to modify registry settings.

2. In the **Component Services** interface, go to **Component Services > Computers**. Right-click **My Computer** and select **Properties**.
3. In the **My Computer properties** panel, go to **COM Security > Access Permissions: Edit Default**.
4. In the **Access Permission** panel, add the Network Service account with Local Access permission set to Allow.
5. Apply your changes and close all open **Component Services** panels.
6. Restart the VM guest.
7. Back up the VM guest and verify that the event ID 8194, VSS error message is not issued to the event log.

Data Protection for VMware installation failure: deployment engine initialization

The Data Protection for VMware installation might be stopped due to a deployment engine initialization failure due to .lock files. If the deployment engine interferes with the Data Protection for VMware installation, the following error message is produced:

```
Deployment Engine failed to initialize.  
The installer will now shutdown. Please check with the log files for a more  
complete description of the failure.  
PRESS ENTER TO CONTINUE:
```

The cause might be deployment engine .lock files that come from a concurrent installation that is running or from an installation that stopped before it completed. If another installation is running, wait until that installation finishes before you install Data Protection for VMware. If there are no other installations that are started and you encounter this problem, delete any .lock files.

Important: Do not delete any .lock files if there are other Data Protection for VMware installations running.

Windows To delete .lock files on Windows, issue the following command:

```
cd C:\Program Files\IBM\Common\acsi\logs
del .lock*
```

Linux To delete .lock files on Linux, issue the following command:

```
cd /usr/ibm/common/acsi/logs
rm .lock*
```

After you remove these files, restart the installation.

Unsupported characters in VM or datacenter name

Data Protection for VMware does not support backing up VMs or datacenters that contain any of the following characters in their name:

- " Double quotation mark
- ' Single quotation mark
- :
- Colon
- ;
- Semicolon
- *
- Asterisk
- ?
- Question mark
- ,
- Comma
- <
- Less than sign
- >
- Greater than sign
- /
- Forward slash
- \
- Backward slash
- |
- Vertical bar

Issues encountered after changing the vCenter

After you change the vCenter in the Data Protection for VMware vSphere GUI, the following two issues might occur:

- A data center that is associated with the new vCenter does not appear on the **Configuration Status** page.

To resolve this issue, manually set the domain for the new vCenter. See [“Set_domain” on page 108](#) for details about issuing this command.

- In the Restore tab, an ESX host (associated with a previous data center) displays under a new data center within the new vCenter. This issue is a known limitation. There is not a workaround for this issue.

Consolidating VM backups

After a VM backup, the VM might contain preexisting snapshots even though no snapshots are present in the Snapshot Manager. For example, the VM hard disk points to snapshot VMDK files (for example *-000001.vmdk) instead of regular VMDK files. Although preexisting snapshots might be intentionally retained, Data Protection for VMware does not provide a mechanism to verify whether the VMDK points to a valid snapshot. When snapshots are not consolidated, and a VM with preexisting snapshot files is backed up, Data Protection for VMware might report an incorrect size for the backup on the IBM Storage Protect server. Snapshot consolidation also prevents other VMware related issues. As a result, consolidate your snapshots whenever this situation occurs.

To resolve this potential problem, VMware vSphere Client 5.x (or later) notifies you when a VM requires snapshot consolidation. For detailed information, see the following article: [Consolidate Snapshots](#)

For vSphere 4.1 (or earlier), no consolidation option is available. However, the equivalent task is to create a snapshot and then complete a Delete All action.

An error occurred while making the Web server request (GVM0103E)

In this situation, the Data Protection for VMware vSphere GUI shows the following error when you tried to access it:

GVM0103E: An error occurred while making the Web server request. If this error persists, check the network connection with the Web server and verify that the Web server is running.

To prevent this error, make sure that the following conditions exist before you start the GUI:

- **Windows | Linux** The Data Protection for VMware vSphere GUI is installed on a system that meets the operating system prerequisites. It must have network connectivity to the following systems:
 - vStorage Backup Server
 - IBM Storage Protect server
 - vCenter Server (Data Protection for VMware vSphere GUI)
- **Windows** The Data Protection for VMware vSphere GUI host URL address must be set in your Internet Explorer trusted sites zone. In the Internet Explorer menu bar, go to **Tools > Internet Options > Security > Trusted sites**. Click **Sites** and add the host URL address. Make sure to apply your changes. For example:

Add this website to the zone:http://myvctrmachine.xyzco.com

Return codes for VM backup operations

The following return codes apply to VM backup operations on Linux or Windows systems:

Table 11. Return codes for VM backup operations	
Return Code	Description
0	A command to process one or more VMs completed successfully.
8	A command to process multiple VMs succeeded on only some of the VMs that were targeted by the command. Examine the log file to determine the processing status for each of the targeted VMs.

Table 11. Return codes for VM backup operations (continued)	
Return Code	Description
12	A command was issued to process one or more VMs. The command did not complete for any of the VMs that were targeted by the command. Examine the log file to determine possible reasons for the failure.

Troubleshooting file restore operations

You can retrieve diagnostic information to resolve file restore issues, by running Microsoft Windows PowerShell cmdlet commands.

Before you begin

Ensure that Microsoft Windows PowerShell 3 or later is available on the system where Data Protection for VMware is installed. To view which version of PowerShell is installed, enter the following command in a PowerShell session:

```
PS C:\> $PSVersionTable.PSVersion
```

The number that displays in the Major column is the PowerShell version.

About this task

Complete these steps on the system where Data Protection for VMware is installed.

Procedure

1. Start a Microsoft Windows PowerShell or Microsoft Windows PowerShell ISE session with administrator authority.

For example:

Start > All Programs > Accessories > Windows PowerShell.

Right-click **Windows PowerShell** and select **Run as administrator**.

2. Verify that execution policy is set to RemoteSigned by issuing the following command:

```
PS C:\> Get-ExecutionPolicy
```

If another policy is shown, set the execution policy to RemoteSigned by issuing the following command:

```
PS C:\> Set-ExecutionPolicy RemoteSigned
```

This policy allows the `vetools.psm1` script to run on the system.

Tip: The **Set-ExecutionPolicy** command must be issued only once.

3. Import the Data Protection for VMware PowerShell module to make the cmdlets available:

```
PS C:\> Import-Module C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\vetools.psm1
```

4. Display log file information in a PowerShell Viewer by issuing the following command:

```
PS C:\> Show-VeFileRestoreLogEntries
```

You can investigate and share log information in the PowerShell Viewer with any of the following actions:

- Enter a term to filter the results.
- Click **Add criteria** to filter the information by more detailed specifications.
- Click one or more rows to save or copy their content for sharing.

5. Optional: Display trace information from a trace file by issuing the following command:

```
PS C:\> Show-VeFileRestoreTraceEntries
```

6. Optional: If you need to gather logs to review detailed diagnostic information (-review) or to send to IBM Support, save the logs in a compressed file by issuing the following command:

```
PS C:\> Get-VeProblemDeterminationInfo -review
```

By default, this command saves the `VeProblemDetermination.zip` file on the desktop.

Tip: If this command returns an error in the default "PowerShell" interface, start the "PowerShell ISE" interface as an administrator. Then, run the command again.

7. Optional: Each cmdlet provides parameters.
To view parameters, issue the following **help** command:

```
help cmdlet name -ShowWindow
```

Related information

[File restore log activity options](#)

[Log file activity](#)

Trace options for file restore

By setting tracing options in the `FRLog.config` file, you can troubleshoot problems that you might encounter during file restore operations.

Modify the options in the `FRLog.config` file with a text editor in administrator mode. The `FRLog.config` file is in the following directory:

```
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\logs
```

FR.API.TRACE=ON | OFF

Specify whether to trace API activity at the recommended level of detail.

Note: The following values are also supported and indicate the least, recommended, and highest level of detail: `DEBUG`, `TRACE`, `ALL`.

API_MAX_TRACE_FILES=number

Specify the maximum number of trace files to be created or used. The default value is 8.

API_MAX_TRACE_FILE_SIZE=number

Specify the maximum size of each trace file in KB. The default value is 8192 KB.

API_TRACE_FILE_NAME=API_trace_file_name

Specify the name of the API trace file. The default value is `fr_api.trace`.

API_TRACE_FILE_LOCATION=API_trace_file_location

Specify the location of the API trace file. Specify the location by using a forward slash (/).

The default location is `Install_Directory/IBM/tivoli/tsm/tdpvmware/webserver/usr/servers/veProfile/logs`.

File restore solutions

Resolve unique or infrequent issues that interfere with file restore operations.

Log in issues

In this scenario, the following information message displays when a fully qualified host name (`myhost.mycompany.com`) or numeric IP address (`192.0.2.0`) is entered in the login page:

```
The host cannot be found. Verify the host name and log in again.  
If the problem persists, contact your administrator.
```

To resolve this issue, enter either the fully qualified domain name (myhost.mydomain) or the short host name (myhost).

File restore operations limitations

Consider the following known issues and limitations for file restore operations issued from the IBM Storage Protect file restore interface while troubleshooting.

Platform specific issues

You may also consult the following topics for issues particular to your platform:

- For Linux issues, see the topic [Linux file restore operations limitations](#)
- For Windows issues, see the topic [Windows file restore operations limitations](#)

The following areas address limitations related to generic interface operations.

Object number limitations

A maximum of 2,000 objects can be restored from search results in a single operation. To restore more than 2,000 objects from search results in the file restore user interface, you must run multiple restore operations.

File restore interface cannot mount a renamed virtual machine

In this scenario, a virtual machine is backed up to the IBM Storage Protect for Virtual Environments server by using the full VM incremental-forever backup type. After the backup completes, the virtual machine is renamed in the VMware vCenter Inventory. A subsequent attempt to mount the virtual machine backup displays the following error message in the file restore interface:

The system cannot find a backup to load. Contact your administrator and log out.

The backup cannot be loaded because the IBM Storage Protect for Virtual Environments server does not contain a backup image of the renamed virtual machine.

To address this error, apply either of the following solutions:

- Interim solution: in the VMware vCenter Inventory, rename the virtual machine to the name that was used in the original backup. Then log in to the file restore interface to access the backup for a file restore operation.
- Long-term solution: back up the renamed virtual machine by using the full VM incremental-forever backup type. Then log in to the file restore interface to access the backup for a file restore operation. The original virtual machine and the renamed virtual machine both appear in IBM Storage Protect for Virtual Environments server database and storage pool. To prevent any confusion, the IBM Storage Protect for Virtual Environments administrator must delete the file space for the original virtual machine.

Log in attempt with host name or IP address fails

In this scenario, the correct host name or IP address is entered in the login page of the file restore interface. However, the following error message is displayed:

The host cannot be found. Verify the host name and log in again. If the problem persists, contact your administrator.

This scenario occurs when the data mover system and the VMware vCenter Server use different internet protocols. For example, the data mover uses IPv4 and the vCenter Server uses IPv6. For example, the data mover uses IPv4 and the vCenter Server uses IPv6. To avoid a login failure, complete either of the following tasks:

- Enter the fully qualified domain name in the login page of the file restore interface. For example: `myhost.mycompany.com`.
- If IPv4 is not used in your environment, request the domain administrator to remove any IPv4 entries from the domain DNS server for that host name. If IPv6 is not used in your environment, request the domain administrator to remove any IPv6 entries from the domain DNS server for that host name.

Limitations related to interface operations

In this scenario, the correct host name or IP address is entered in the login page of the file restore interface. However, the following error message is displayed:

```
The host cannot be found. Verify the host name and log in again.
If the problem persists, contact your administrator.
```

"

This scenario occurs when the data mover system and the VMware vCenter Server use different internet protocols. For example, the data mover uses IPv4 and the vCenter Server uses IPv6. To avoid a login failure, complete either of the following tasks:

- Enter the fully qualified domain name in the login page of the file restore interface. For example: `myhost.mycompany.com`
- If IPv4 is not used in the environment, request the domain administrator to remove any IPv4 entries from the domain DNS server for that host name. If IPv6 is not used in the environment, request the domain administrator to remove any IPv6 entries from the domain DNS server for that host name.

Log in can take a long time

Logging in to the file restore interface might take a long time depending on the number of guests that are managed by the VMware vCenter server. For example, it can take three minutes to log on for a vSphere environment with 3,000 guests.

Date stamp, time stamp, and version number do not display in Details or Job History views

During a file restore operation, when a file with the same name exists, the restored file's original modification date and time is appended to the file name. Subsequent restores of the same file contain a version number (_N) after the original modification date and time. For example: `t2.2015-03-07-07-28-03_1.txt`

In this scenario, during a restore operation of a file with the same name as an existing file, the restore operation either fails or is canceled by the user. When viewing information about the failure or cancellation in the **Details** or **Job History** view, the file's date stamp, time stamp, and version number do not display. The original file name displays.

To view the most recently restored version of the file on the guest virtual machine, look for the file with the highest version number.

IBM Storage Protect for Virtual Environments file restore login page requires repeat credentials

This scenario occurs when either the login page is inactive for an extended period of time or the client acceptor service is restarted on the data mover system. After the correct credentials are entered, you are prompted to login again, even though the correct credentials were already entered. To resolve this problem, reenter the correct credentials. The interface then loads the backup selection page.

To avoid this issue, refresh the login page after the client acceptor service is restarted on the data mover or the login page has become inactive.

No subtitles are shown in the file restore product videos

Subtitles (captions) are not available in the file restore product videos at this time. To view the file restore product videos with subtitles, go to the following websites:

- [Finding and restoring files](#)
- [Monitoring file restore operations](#)

Click the **CC** button in the video player to display the subtitles in the videos.

No sound in the file restore product videos on a remote computer

If you view the product videos from the file restore interface on a remote computer that does not have a sound card, you will not hear any sound in the videos. A sound card must be installed and enabled on the remote computer in order for you to hear sound in the product videos.

To view the videos with sound, access the file restore interface from a web browser on a computer that has a sound card installed.

Linux file restore operations limitations

Consider the following known Linux platform issues and limitations for file restore operations issued from the IBM Storage Protect file restore interface while troubleshooting.

Support volume types

The Linux mount proxy system can mount and support volume types such as EXT2, EXT3, EXT4, REISERFS, XFS, BTRFS, RAID device, and LVM device. Using volume types that are not supported by the Linux mount proxy system results in failure of mount operation and the following error message is displayed:

```
ANS3141W Mounting the following device failed: Volume not supported.
```

Directories that are created (or recreated) during a restore operation might be assigned incorrect access permissions, ownership information, or both

Upon a successful restore operation, the original user access permissions and ownership information of the restored files is preserved. However, access permissions and ownership information of the recreated directories (if any) change in accordance with the default umask setting and the initial login group of the user that is logged in for restore on the target Linux system. This is a known limitation.

Linux temporary file system supports file restore to alternate location only

A Linux file system that is mounted in the guest virtual machine, but is not present in the `/etc/fstab` system configuration file, supports a file restore to alternate location operation only. It does not support a file restore operation to the original location. In this scenario, it does not matter whether the file system is listed in the `/etc/mntab` file.

A Linux file system that is mounted in the guest virtual machine, but is not present in the `/etc/fstab` system configuration file, is represented in the file restore user interface as `volume#`. This is a known limitation.

In this scenario, a Red Hat Enterprise Linux (RHEL) 6.5 mount proxy system attempts to mount volumes from an RHEL 7.1, SLES 11, or SLES 12 virtual machine. The operation fails and the mount proxy system reboots. This scenario occurs because a Linux mount proxy system cannot mount Btrfs or XFS volumes from a virtual machine with a later Linux operating system. This is a known limitation.

To prevent this situation, ensure that the minimum level of the Linux mount proxy system is at the same level, or later level, than the level of the protected guest virtual machine.

Linux mount proxy system cannot mount volumes from a virtual machine with a later Linux operating system

In this scenario, a Red Hat Enterprise Linux (RHEL) 6.5 mount proxy system attempts to mount volumes from an RHEL 7.1, SLES 11, or SLES 12 virtual machine. The operation fails and the mount proxy system reboots. This scenario occurs because a Linux mount proxy system cannot mount Btrfs or XFS volumes from a virtual machine with a later Linux operating system. This is a known limitation.

To prevent this situation, ensure that the minimum level of the Linux mount proxy system is at the same level, or later level, than the level of the protected guest virtual machine.

Group identifier is not preserved during restore operation by a non-root user

When a non-root user restores a file that is owned by the same non-root user, the group identifier (GID) for this file is not preserved. This is a known limitation.

To preserve the GID for a file that is owned by a non-root user, restore the file by using root user authority.

Potential UUID collisions on Linux mount proxy system and guest virtual machines

File restore operations are not supported when UUID collisions occur. A UUID collision is defined as the mount proxy system has an identical UUID as the guest virtual machine being whose disks are being mounted. UUID collisions might occur between the guest virtual machines and the mount proxy system in any of the following scenarios:

- The guest virtual machine is cloned from the mount proxy system.
- The mount proxy system is cloned from the guest virtual machine.
- The guest virtual machine and mount proxy system are cloned from the same template.
- The guest virtual machine is cloned from a virtual machine that is already mounted on the mount proxy system.

When a UUID collision exists between the guest virtual machine and the mount proxy system, a file restore operation cannot identify the original restore points on the mounted guest virtual machine. In this situation, the volumes related to the duplicated UUIDs are skipped. Before using the file restore interface, be sure to resolve any UUID collisions between the mount proxy system and guest virtual machines in Linux. Engage your Linux administrator, or Linux support, for guidance.

The following example shows how to generate a new UUID in most file systems, but not all file systems. The workaround is not valid for Btrfs file systems. Ensure your Linux administrator, or Linux support, validates the steps necessary to generate a new UUID for the file systems in use on either the guest virtual machine or mount proxy system.

1. Generate a new UUID on the cloned system by running the `uuidgen` command.
2. Assign the new UUID to a specified block device on the cloned system by running the `tune2fs` command with the `-U` option (`tune2fs <device> -U <new_UUID>`).

LVM devices are not mounted if the LVM importing tool is not able to generate a new UUID

On some Linux distributions, it has been observed that the `vgimportclone` command is not able to assign a new UUID to the device. This results in duplicated device UUIDs on the mount proxy machine. In this situation, the volumes related to the duplicated UUIDs are skipped and the ANS3184W message is reported in the TSM log file. This could be a signal that the guest virtual machine is a clone of the mount proxy system. See also the above section, **Potential UUID collisions on Linux mount proxy system and guest virtual machines**.

Linux Mount Proxy lvm2d daemon may activate the LVM volumes before it is activated by the IBM Storage Protect for Virtual Environments services

This behavior causes the file level restore process not to mount one or more guest filesystem(s) without the mount process issuing any warning or error message. To avoid that, the steps to disable it on the Linux Mount Proxy host are as follows:

1. Edit the `/etc/lvm/lvm.conf` file to set the parameter `use_lvm2d` to 0.
2. Run the following administration commands:
 - `systemctl stop lvm2-lvm2d`
 - `systemctl disable lvm2-lvm2d`
 - `systemctl stop lvm2-lvm2d.socket`
 - `systemctl disable lvm2-lvm2d.socket`

Linux file restore operations limited to a maximum overall path length of 4096 characters

The overall path for the mount in the file restore interface is composed of the following objects:

```
/tsmmount/FLR/<VMName>/<SNAPSHOT_DATE>/VolumeX/dir1/dir2/dir3/dir4/dir5/fileName
```

The maximum path length for the mount in the file restore interface is determined by the total number of combined characters for each object that compose the path of the mount:

- `/tsmmount/FLR/<VMName>/<SNAPSHOT_DATE>/VolumeX`: Maximum length of the mount path: 38 characters.
- `<VMName>`: Maximum length of the virtual machine name: 80 characters.
- `fileName`: Maximum length of the file name to be restored: 255 characters.
- `dir1/dir2/dir3/dir4/dir5`: Maximum length of the path to the file to be restored is determined by the following formula:
 - 4096
 - - length of mount path (38)
 - - length of virtual machine name (80)
 - - length of file name to be restored (255)
 - = Maximum length of the path to the file to be restored

If the total number of combined characters for the path of the mount exceeds 4096 characters, the file is skipped and the restore operation continues for other selected files and directories. No message displays to identify files that are skipped.

If a Linux virtual machine is included in file restore operations, make sure that the total number of combined characters for the path of the mount does not exceed 4096 characters.

Directory does not restore to alternate location on Linux virtual machine

In this scenario, a directory is selected for a Restore to > Alternate Location operation. After the operation completes, the following error message displays for the directory:

```
You do not have permission to restore the file to the destination folder.
```

This error occurs when either of the following conditions exist:

- The directory has the same name as a file that was selected for restore.
- The directory has the same name as a file that exists on the target virtual machine.

This is a known limitation.

SUSE Linux Enterprise Server 12 Btrfs subvolumes are not supported

By default, SUSE Linux Enterprise Server 12 uses Btrfs subvolumes for the root partition. During a file restore operation, the default subvolume is mounted as the root partition. However, Btrfs subvolumes such as /home, /usr, or /opt, are excluded from the snapshot and do not display in the file restore interface. As a result, Btrfs subvolumes are not supported for file restore operations.

A deleted Linux file system restores to the root directory

In this file restore scenario, a Linux file system (on the guest virtual machine) is unmounted and the file system mount point is deleted. A subsequent attempt to restore the file system restores the file system data to the root (/) file system. This issue occurs because the file system to restore is unmounted. As a result, only the mount point is recreated, and the data is restored to that directory under the root (/) file system.

Linux SLES 11 mount proxy system cannot mount Btrfs or XFS file systems from a Linux SLES 12 guest virtual machine

This environment consists of a SUSE Enterprise Linux 11 (Service Pack 3) mount proxy system, and a SUSE Enterprise Linux 12 guest virtual machine. In this scenario, the SUSE Enterprise Linux 11 mount proxy system cannot mount files from a SUSE Enterprise Linux 12 Btrfs or XFS file system. The mount operation fails and the following message displays:

```
An error occurred while loading the backup ...
```

Local user authentication is required for a virtual machine

Make sure that the user who authenticates to the Linux virtual machine that contains the files to be restored is a local user. Authentication is not available through Windows domain, Lightweight Directory Access Protocol (LDAP), Kerberos, or other network authentication methods.

File restore operations from a local VVOL snapshot require a free loop device for each of the guest virtual machine

Old Linux distributions have a limited number of default loop devices and this can limit the number of virtual machines that can be mounted on the mount proxy system. For some distributions, this number is set to 8 or to 64. This number can be theoretically set to a maximum of 256 loop devices, but requires a reboot of the machine and for some Linux distributions, to recompile part of the kernel.

Installing the Linux mount proxy on old Linux distributions (for example, Suse 11) can limit the number of disks that can be mounted on the mount proxy for file restore VVOL operations. In this situation, the file restore operation is stopped and a cleanup is performed.

For example, if a machine has 8 loop devices and only 6 are free, it is not possible to mount a local snapshot of a machine with 7 or more disks.

Use a Linux distribution such as Suse 12 or Redhat 7.3 for the mount proxy. These distributions dynamically create the loop devices upon request.

File restore operations from a local VVOL snapshot create a temporary defunct dsmagent on the Linux mount proxy system

For each mount of a snapshot of a VMDK, a defunct dsmagent is displayed in the process list of the mount proxy. This is due to incorrect handling inside the VMware library. The defunct dsmagent is no longer displayed after the main dsmagent has stopped.

Windows file restore operations limitations

While troubleshooting, consider the following known Windows platform issues and limitations for file restore operations issued from the IBM Storage Protect file restore interface.

Directory does not restore to alternate location on Windows virtual machine

In this scenario, a directory is selected for a **Restore to > Alternate Location** operation. After the operation completes, the following error message displays for the directory:

```
File not found during backup, Archive or Migrate processing.  
No file specification entered.
```

This error occurs when either of the following conditions exist:

- The directory has the same name as a file that was selected for restore, and the file is restored before the directory.
- The directory has the same name as a file that exists on the target virtual machine.

This is a known limitation.

Microsoft Windows user account cannot access its own user folders

User folders are a collection of folders that are created by the Windows operating system for each user account. These folders are typically in `C:\Users\<user account>`. A Windows user account cannot access its own user folders in the file restore interface. This is a known limitation.

Directories that are restored to an alternate location by a Windows domain administrator become owned by the domain administrator

In this scenario, a Windows domain administrator restores a directory and files that are owned by a Windows domain user, and **Restore to > Alternate Location** is selected. The restore operation completes successfully, and the Windows domain user retains ownership of the restored files. However, ownership of the restored directory changes to the Windows domain administrator. This is a known limitation. (

Windows operating system identifies date and time stamp as a file type

During a file restore operation, when a file with the same name exists, the restored file's original modification date and time is added to the file name. For example, the file `1.txt` can contain this date and time stamp: `1.2015-5-10-04-00.txt`

However, a Windows operating system identifies the date and time stamp as a file type. For example, when you view the file `1.2015-5-10-04-00.txt` in Windows Explorer, the **Type** column displays `2015-5-10-04-00.txt`. The file format is not affected. This is a known limitation.

VMware attributes

Review how Data Protection for VMware interacts with VMware attributes.

VMware custom attributes

Data Protection for VMware includes VMware custom attributes in backup and restore operations. However, custom attributes are only included when the data mover node is connected directly to a vCenter Server and not an ESXi Server. To set this connection, specify a vCenter Server with the `VMCHost` option that is on the data mover node.

For more information, see the following VMware Knowledge Base article:
<http://kb.vmware.com/kb/1005720>

VMware configuration attributes

Data Protection for VMware does not directly access, modify, or back up the .vmx file. The .vmx file is non-portable. As a result, Data Protection for VMware does not save values that are contained in the file or save the actual file. The main objective of Data Protection for VMware is to recover the VM to a usable (or startable) state.

Accepting self-signed SSL certificates

You are prompted to accept security certificates the first time you connect to or monitor remote data movers.

Each data mover installation runs on a web service that creates and maintains a self-signed Secure Sockets Layer (SSL) certificate. The certificate secures interactions with the data mover, such as establishing connection, monitoring or upgrading. When the IBM Data Protection for VMware vSphere plug-in initiates communication with each remote data mover, you are prompted to accept a security certificate. After the certificate is accepted, that information is retained for future communications with the same remote data mover. You are presented with the details of the certificate the first time a connection is established. You can accept or reject the certificate.

To help ensure the security of the connection, you can compare the contents of the self-signed certificate presented in a browser by IBM Data Protection for VMware in the vSphere plug-in with the equivalent certificate on the installed web server. To do this, set up a keytool to access the keystore. For instructions, see [“Verifying a Data Protection for VMware self-signed web server certificate”](#) on page 177.

If you accept the self-signed certificate, subsequent connections between the host and data mover are digitally signed and will not be challenged. If the local file is deleted, a challenge to accept a new digital certificate will be issued on the next operation.

If you reject the self-signed certificate, the connection is refused.

Verifying a Data Protection for VMware self-signed web server certificate

For enhanced security, you can verify the contents of a web server Secure Sockets Layer (SSL) self-signed certificate produced by Data Protection for VMware against the equivalent certificate on the installed certification server.

About this task

When you are presented with the details of a web server SSL certificate the first time a connection is established in the plug-in, you can accept or reject the certificate. At this point, you might want to verify the contents of the certificate that you have received with the equivalent certificate in the web server keystore. The following instructions are for a Liberty keystore. For other web server-based certificate services, see the system documentation.

Procedure

To access and verify certificate details in the keystore:

1. On the workstation on which the Data Protection for VMware vSphere GUI is installed, ensure that the JAVA_HOME environment variable is correctly set. Complete the appropriate actions for your operating system:
 - a) Obtain the Java virtual machine version by going to the C:\Program Files\Common Files\Tivoli\TSM directory and noting the value of the subdirectory jvm. For example, if the directory name is "jvm80520", note the numeric value 80520.
 - b) To set the environment variable, issue the following command: `set JAVA_HOME=C:\Program Files\Common Files\Tivoli\TSM\jvm80520\jre.`
 - a) Run the following command:

```
export JAVA_HOME=/opt/tivoli/tsm/tdpvmware/common/jre/jre
```

2. Add the keytool to your path:

Windows

```
set PATH=%JAVA_HOME%\bin;%PATH%
```

Linux

```
export PATH=$JAVA_HOME/bin:$PATH
```

If the JAVA_HOME and PATH variables are correctly set, you can now invoke the keytool.

3. At the command prompt, enter **keytool** and press **Enter**.

If you receive the **keytool** help output, the tool is correctly configured. If the Command not found message is displayed, verify that the JAVA_HOME and PATH variables were updated correctly.

4. Go to the location of the key store:

Windows

```
cd C:\IBM\StorageProtect\webserver\usr\servers\veProfile\resources\security
```

Linux

```
cd /opt/tivoli/tsm/tdpvmware/common/webserver/usr/servers/veProfile/
resources/security
```

5. At the command prompt, list the directory contents. Verify that the certificate file key.jks is present in the directory listing.
6. Issue the following command:

```
keytool -list -keystore key.jks -v
```

7. When prompted for the password, press **Enter**.

The password is randomly generated and enables only the ability to verify the certificate details that the prompt presents.

8. Compare the output from the keytool with that of the certificate prompt. The serial number and fingerprints of the self-signed certificate and the equivalent certificate on the keystore must match.
9. If the certificate details do not match, contact the administrator of the web server-based certificate services.

Troubleshooting IBM Storage Protect vSphere Client plug-in problems

Solutions are provided for IBM Storage Protect vSphere Client plug-in issues. You can learn how to resolve Platform Services Controller connection problems, enable tracing, and get more details about IBM Storage Protect vSphere Client plug-in messages.

- [“Resolving Platform Services Controller connection problems” on page 178](#)
- [“Enabling tracing” on page 179](#)
- [“Messages for the IBM Storage Protect vSphere Client plug-in” on page 181](#)

Resolving Platform Services Controller connection problems

Tags and categories that are used for the management of virtual machine backups are stored and managed on the VMware Platform Services Controller (PSC). To be able to use the tagging feature for data protection, the tag-based data mover node and the IBM Storage Protect vSphere Client plug-in must be able to connect to the Platform Services Controller by using the Single Sign On process.

The Platform Services Controller server hosts the VMware Lookup Service that registers the location of vSphere components and handles the vCenter Single Sign On process.

Symptoms

When connection problems occur, the data mover node cannot complete the Single Sign On process and cannot access the tags and categories in the Platform Services Controller.

If the Platform Services Controller cannot be reached, the tag information will not be displayed in the IBM Storage Protect vSphere Client plug-in. Virtual machine backup operations will also fail.

Resolving the problem

Complete the following tasks to diagnose and resolve connectivity problems:

- Ensure that the Platform Services Controller host is powered on and accessible over the network.
- Ensure that the VMware Lookup Service is active and accepting connections at the following address: `https://PSC-FQDN/lookupservice/sdk`, where *PSC-FQDN* is the fully qualified domain name of the Platform Services Controller host.
- Ensure that a data mover is installed on the same server that hosts the Data Protection for VMware vSphere GUI. The data mover node must be configured so that the vCenter server credentials are saved, for example, by using the **dsmc set password** command in the backup-archive command-line.
- On UNIX and Linux clients, the existing passwords in the `TSM.PWD` files are migrated to the new password store in the same location. For root users, the default location for the password store is `/etc/adsm`. For non-root users, the location of the password store is specified by the `passworddir` option.

The `TSM.PWD` file is deleted after the migration.

- Ensure that client option `vmhost` is set by using the same value and format that was used for the vCenter server field during the installation of Data Protection for VMware. The preferred format for the vCenter server address is the vCenter server's fully qualified domain name (FQDN). Use the vCenter server IP address only if it was used during the registration of the vCenter, although the IP address is not preferred by VMware.
- The system time on the data mover host must be in sync with the system time on the Platform Services Controller and vCenter. The system time and time zone must be set correctly on all three systems. Otherwise, a Platform Services Controller connection error occurs. The following message is typical of this type of error:

```
ANS2378E Single Sign On login to the vSphere Server failed in function  
visdkGetSecurityToken - Issue. "The time now Wed Apr 20 21:31:58 UTC 2016  
does not fall in the request lifetime interval extended with clock tolerance  
of 600000 ms: [ Wed Apr 20 16:20:46 UTC 2016; Wed Apr 20 16:50:46 UTC 2016).  
This might be due to a clock skew problem."
```

- For more information about messages that occurred, see [“Messages for the IBM Storage Protect vSphere Client plug-in” on page 181](#).

Enabling tracing

By enabling the tracing feature, you can troubleshoot problems that you might encounter during operations with the IBM Storage Protect vSphere Client plug-in or the tag-based data mover node.

About this task

To enable tracing in the common VMware vCloud Suite layer for both the backup-archive command-line client and the IBM Storage Protect vSphere Client plug-in, the following trace files and trace properties files are used:

Log location

Trace output is added to the following log files:

- **Windows** (Data Protection for VMware)
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\logs\vcspplugin.log
- **Linux** (Data Protection for VMware) /opt/tivoli/tsm/tdpvmware/common/webserver/usr/servers/veProfile/logs/vcspplugin.log

Log property location

The property values are updated in the following trace properties files to enable tracing:

- **Windows** (Client) C:\Program Files\Tivoli\TSM\baclient\plugins\vcloudsuite\sdk\log4j.properties
- **Linux** (Client) /opt/tivoli/tsm/client/ba/bin/plugins/vcloudsuite/sdk/log4j.properties
- **Windows** (Data Protection for VMware)
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\apps\tsmVmGUI.war\WEB-INF\classes\log4j.properties
- **Linux** (Data Protection for VMware) /opt/tivoli/tsm/tdpvmware/common/webserver/usr/servers/veProfile/apps/tsmVmGUI.war/WEB-INF/classes/log4j.properties

Procedure

1. To view more detailed trace information for the common VMware vCloud Suite layer, change the following property value from INFO to TRACE in the corresponding log4j.properties file:

```
logger.ve.level = TRACE
```

2. To view more detailed trace information for the common vCloud Suite layer, VMware vCloud Suite SDK, and associated .jar files, change the following property value from WARN to TRACE, in the corresponding log4j.properties file

```
rootLogger.level =TRACE
```

3. Rerun the actions or commands that caused the error.

If the error occurred in the IBM Storage Protect vSphere Client plug-in, you must restart the server that hosts the Data Protection for VMware vSphere GUI.

Resolving administrator ID not found messages

Some data and options are not shown in IBM Storage Protect vSphere Client plug-in if the IBM Storage Protect server administrator ID is not available. This ID must be configured in the Data Protection for VMware vSphere GUI.

Procedure

If data or options are not shown in IBM Storage Protect vSphere Client plug-in and a message is displayed specifying that the administrator ID cannot be found, complete the following steps in the Data Protection for VMware vSphere GUI to set and save the administrator ID:

1. In the GUI menu bar, click **Configuration**.
2. Select **Edit IBM Storage Protect Configuration** in the **Tasks** menu.
3. On the **Server Credentials** page, complete the following steps:
 - a) Enter the administrator ID in the **IBM Storage Protect Admin ID** field if it is not already entered and complete the corresponding fields and options for the password and port.
 - b) Select the **Save the administrator ID, password, and port settings for use in future sessions** check box.

If an administrator ID is configured, but this check box is not selected, the administrator ID will not be available for IBM Storage Protect vSphere Client plug-in sessions.

4. Click **OK** to save the changes.

Messages for the IBM Storage Protect vSphere Client plug-in

To help you understand IBM Storage Protect vSphere Client plug-in messages, review the following information:

- [GVM5107E](#)
- [GVM5110E](#)
- [GVM5111E](#)
- [GVM5112E](#)

GVM5107E: Data protection settings are not available because the login credentials provided are invalid for the 'name' Platform Services Controller

The symptoms, cause, and user response are provided for this IBM Storage Protect vSphere Client plug-in message.

Symptoms

Data protection settings cannot be displayed in the IBM Storage Protect vSphere Client plug-in.

Causes

The credentials that are required to log in to the Platform Services Controller are invalid for the vCenter.

Resolving the problem

Ask the IBM Storage Protect server administrator to update the vCenter Server credentials by using the **dsmc set password** command in the backup-archive command-line client on the server where the Data Protection for VMware vSphere GUI is installed.

```
dsmc set password -type=vm vmhost username password
```

The value for the `vmhost` option must match the value that is in the client options file. It must also match the vCenter server address that was used during the installation of the Data Protection for VMware vSphere GUI.

You might also receive the following data mover messages that are associated with this error:

- [Client message ANS9331W](#)
- [Client message ANS9332E](#)

Related information

[Set Password](#)

GVM5110E: Data protection settings are not available because an error occurred connecting to the 'name' Platform Services Controller

The symptoms, cause, and user response are provided for this IBM Storage Protect vSphere Client plug-in message.

Symptoms

Data protection settings cannot be displayed in the IBM Storage Protect vSphere Client plug-in.

Causes

Other errors are causing connection issues to the Platform Services Controller and vCenter.

Resolving the problem

Ensure that the Platform Services Controller listed is running as expected. For more information, see [“Resolving Platform Services Controller connection problems”](#) on page 178.

You might also receive the following data mover message that is associated with this error:

- [Client message ANS2373E](#)

GVM5111E: Data protection settings are not available because no login credentials are found for the '*name*' Platform Services Controller

The symptoms, cause, and user response are provided for this IBM Storage Protect vSphere Client plug-in message.

Symptoms

Data protection settings cannot be displayed in the IBM Storage Protect vSphere Client plug-in.

Causes

The credentials that are required to log in to the Platform Services Controller cannot be found for the vCenter. The credentials might not be found because they were never stored, the permissions to the credentials are not valid or incorrect credentials were stored.

Resolving the problem

Windows

Ask the IBM Storage Protect server administrator to update the vCenter Server credentials by using the **dsmc set password** command in the backup-archive command line client on the server where the Data Protection for VMware vSphere GUI is installed.

```
dsmc set password -type=vm vmhost username password
```

The value for the `vmhost` option must match the value that is in the client options file. It must also match the vCenter server address that was used during the installation of the Data Protection for VMware vSphere GUI.

After the **dsmc set password** command runs, the credentials are stored again.

Linux

If you suspect a permissions issue for the credentials, complete the follow steps:

1. Ensure that the `TSM.IDX`, `TSM.KDB`, and `TSM.sth` files are accessible by the `tdpvmware` user.
2. If you receive a permission error, check that the permissions for the above three files appear as follows:

```
-rw-r----- 1 root tdpvmware
```

If the permissions do not match, go to the directory that contains the above three files (`/opt/tivoli/tsm/tdpvmware/common/webserver/usr/servers/veProfile/tsmVmGUI`) and change the permissions for each of these files with the following commands, for example:

```
chgrp tdpvmware TSM.IDX  
chmod g+r TSM.IDX
```

You might also receive the following data mover messages that are associated with this error:

- [Client message ANS9331W](#)
- [Client message ANS9332E](#)

Related information

Set Password

GVM5112E: Data protection settings are not available because an error occurred processing the login credentials that are provided for the '*name*' Platform Services Controller

The symptoms, cause, and user response are provided for this IBM Storage Protect vSphere Client plug-in message.

Symptoms

Data protection settings cannot be displayed in the IBM Storage Protect vSphere Client plug-in.

Causes

The credentials that are required to log in to the Platform Services Controller are available but one or more of the following errors occurred:

- An error occurred processing the credentials.
- An error occurred loading the native library that is needed to process the credentials.

Resolving the problem

Contact the IBM Storage Protect server administrator for assistance.

You might also receive the following data mover messages that are associated with this error:

- [Client message ANS2635E](#)
- [Client message ANS9365E](#)

Appendix B. IBM Storage Protect recovery agent operations

This service enables the mounting of any snapshot volume from the IBM Storage Protect server. You can use the iSCSI protocol to access the snapshot from a remote computer.

If you need to view the snapshot locally with read-only access on the client system, use Data Protection for VMware 8.1.4 or earlier versions.

In addition, the recovery agent provides the instant restore function. A volume that is used in instant restore processing remains available while the restore process proceeds in the background. The recovery agent is accessed with the recovery agent GUI or command-line interface.

Important: Previous versions of IBM Storage Protect for Virtual Environments delivered mount and file restore function with the recovery agent. Although this function is still supported by the recovery agent, the IBM Storage Protect file restore interface is the preferred method, as described in the following topic:

[Chapter 3, “Getting started with file restore,” on page 51](#)

The content in this "IBM Storage Protect recovery agent operations" collection is provided as a reference for users who prefer the recovery agent method.

Windows Mounting snapshots with the recovery agent

You can use the IBM Storage Protect recovery agent to mount a snapshot and use the snapshot to complete data recovery.

Mount snapshots with either the recovery agent GUI or with the [“Mount command” on page 122](#). Install and run the recovery agent on a system that is connected to the IBM Storage Protect server through a LAN. You cannot use the recovery agent component operations in a LAN-free path.

Be aware of these three situations when running mount operations:

- When the recovery agent is installed on a guest machine, you cannot start an instant restore or a mount operation for any file system or disk while the guest machine 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 VM.
- When you browse the snapshot backup inventory, the operating system version of the VM is the version that was specified when the VM was originally created. As a result, the recovery agent might not reflect the current operating system.
- A volume becomes unstable when a network failure interrupts a mount operation. 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 recovery agent GUI.

A maximum of 20 iSCSI sessions is supported. The same snapshot can be mounted more than one time. If you mount a snapshot from the same tape storage pool by using multiple instances of the recovery agent, one of the following actions occurs:

- The second recovery agent instance is blocked until the first instance is complete.
- The second recovery agent instance might interrupt the activity of the first instance. For example, it might interrupt a file copy process on the first instance.
- The recovery agent cannot connect to multiple servers or nodes simultaneously.

As a result, avoid concurrent recovery agent sessions on the same tape volume.

Mounting snapshot guidelines

Snapshots can be mounted in either read-only or read/write mode. In read/write mode, the recovery agent saves changes to data in memory. If the service is restarted, the changes are lost.

The recovery agent operates in either of the following two modes:

No user is logged in

The recovery agent runs as a service. This service enables remote connections through the Data Protection for VMware command-line interface.

User is logged in

The recovery agent continues to run as a service until you start the recovery agent and use the GUI. When you close the recovery agent and GUI, the service restarts. You can use only the recovery agent application and GUI when running with administrator login credentials. Only one copy of the recovery agent application can be active at any time.

When mounted volumes exist and you start Mount from the Start menu on Microsoft Windows operating systems, 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 unmounted, even when they are in use.

Restriction: When exposing snapshots as iSCSI targets, and a snapshot of a dynamic disk is displayed to its original system, the UUIDs become duplicated. Likewise when a snapshot of a GPT disk is displayed to its original system, the GUIDs become duplicated. To avoid this duplication, expose dynamic disks and GPT disks to a system other than the original system. For example, expose these disk types to a proxy system, unless the original disks no longer exist.

Windows Restoring files with the recovery agent

Use the IBM Storage Protect recovery agent for efficient file restores and to minimize downtime by mounting snapshots to virtual volumes.

You can use the recovery agent for the following tasks:

- Recovering lost or damaged files from a backup
- Mounting a VM guest volume and creating an archive of the VM guest files
- Mounting database applications for batch reports

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, the 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 unmounting the volume.

The default *read only* mount option is the preferred method, unless a mounted volume must be writeable. For example, an archive application might require write access to the archived volume.

The recovery agent mounts snapshots from the IBM Storage Protect server. In the recovery agent GUI, click **Remove** to close an existing connection to a server. You must remove any existing connection before you can establish a new connection to a different server or different node. Dismount all volumes before you click **Remove**. The remove operation fails if there are active mount and restore sessions in the Windows Mount machines. You cannot remove the connection to a server when you are running 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 the recovery agent. Otherwise, these mounted virtual volumes cannot be unmounted after the recovery agent is reinstalled.

Restoring file information for a block-level snapshot is a random-access process. As a result, processing might be slow when a sequential-access device (such as a tape) is used. To run a file restore of data that is stored on tape, you must first move the data to disk or file storage as file restore from tape media is no longer supported. From the IBM Storage Protect server administrative command-line client (dsmadm), issue the **QUERY OCCUPANCY** command to see where the data is stored. Then, issue the **MOVE NOEDATA** command to move the data back to disk or file storage.

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.

Restriction: Do not change the IBM Storage Protect node password while running a file restore or an instant restore from snapshots stored in that node.

Windows | Linux **Restoring volumes instantly with the recovery agent**

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.

Instant restore works only with local volumes. The term "local" is used regarding the IBM Storage Protect recovery agent since it must be installed on the guest machine that contains the volume to be restored. Local volumes must have an assigned drive letter. Instant restore cannot be used to restore the system volume.

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 Software RAID 0, RAID 1, and RAID 5 volumes. You can use a basic disk as a destination volume and then convert the basic disk to a dynamic disk. The file system on the destination volume 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 might not be affected, and you can work with the cluster, and with that volume, in parallel. During the instant restore, the disk that is 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.

Windows Restoring volumes instantly from a Windows system with the recovery agent

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.

Before you begin

Important: Previous versions of IBM Storage Protect for Virtual Environments delivered mount and file restore function with the IBM Storage Protect recovery agent. Although this function is still supported by the recovery agent, the IBM Storage Protect file restore interface is the preferred method, as described in the following topic:

[Chapter 3, “Getting started with file restore,” on page 51](#)

The content in this "IBM Storage Protect recovery agent operations" collection is provided as a reference for users who prefer the recovery agent method.

Instant restore operations on Windows require the recovery agent to be installed on the guest machine.

Instant restore is available only from Data Protection for VMware snapshots on a source volume that is on a simple, MBR-based disk. 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 is 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 that are 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.

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 restarts, the volume appears to be unformatted. Do not attempt to format or modify the volume. 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 IBM Storage Protect 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 into a simple volume that is on a dynamic disk. However, the source volume must be an MBR-based disk. 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**.

Before proceeding, make sure you have reviewed the following information:

- “Mounting snapshots with the recovery agent” on page 185
- “Restoring volumes instantly with the recovery agent” on page 187

The recovery agent GUI must be configured before attempting a file restore operation. To configure, click **Select IBM Storage Protect server** and **Settings** in the recovery agent GUI and enter the required information.

About this task

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 IBM Storage Protect 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**.

Procedure

To perform an instant restore, complete the following steps:

1. On the guest machine, start the recovery agent GUI from the **Start > All Programs** menu or by clicking the recovery agent icon in the taskbar.
2. In the recovery agent window, select the IBM Storage Protect server to use as the source by clicking **Select IBM Storage Protect server**. Although the **Select IBM Storage Protect server** list appears to contain multiple servers, this list contains a maximum of one server only.
The recovery agent queries the server for a list of protected VMs and displays the list.
3. Select a VM, date, time, and disk, and then click **Restore**.
4. The 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).
5. 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.
 6. Click **Restore**.
 7. 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.

Appendix C. Data Protection for VMware vSphere GUI messages

This information contains explanations and suggested actions for messages issued by the Data Protection for VMware vSphere GUI.

For messages shown in the Data Protection for VMware vSphere GUI that contain the FMM prefix, message information is available at the following web site: [FMM, FMF, FMV, FMX, FMY: IBM Storage Protect Snapshot messages](#)

GVM messages

GVM0001E	The operation failed with return code <i>return code</i>
-----------------	---

GVM0002E	An internal error occurred: <i>type of error</i>
-----------------	---

GVM0003E	A connection with the IBM Storage Protect server could not be established.
-----------------	---

Explanation:

The server might not be running.

Administrator response:

Check the network connection with the server machine. Verify that the server is running and try to log in again.

GVM0004W	Are you certain that you want to delete this data?
-----------------	---

Explanation:

You cannot recover the data after it is deleted. Ensure that the data is not needed before you delete it.

Administrator response:

Click OK to delete the data or click Cancel to cancel this action.

GVM0005W	The connection with the IBM Storage Protect server has timed out.
-----------------	--

Explanation:

Possible causes include a long-running operation, a problem on the server, or a communications problem.

Administrator response:

If the operation is long-running, the operation might be complete or it might soon be complete. Before trying the operation again, determine if the expected result occurred. Check the activity log of the IBM Storage Protect server for errors related to the operation. Using a SSL port without selecting SSL can cause this error.

GVM0006I	A server connection with the name <i>server name</i> has been successfully created. Click OK to continue.
-----------------	--

GVM0007W	There is no IBM Storage Protect server definition found.
-----------------	---

Explanation:

A connection for a IBM Storage Protect server must be defined before any server operations or queries are performed.

Administrator response

To define a server:

1. Click the Configuration tab.
2. Click the Edit Configuration Settings action link.
3. Click the IBM Storage Protect Server Credentials tab.

GVM0008E	An error occurred while writing to the server's database file, <i>tsmsrvr.props</i>
-----------------	--

Explanation:

The server definition could not be written to the *tsmsrvr.props* file.

Administrator response:

The file must reside in the install directory of IBM Storage Protect. Before you try the action again, verify that the file exists and that the file is not write protected.

GVM0011I	The VM <i>VM name</i> is spanned into multiple datastores. It can only be restored to its original location.
-----------------	---

GVM0011W	The VM <i>VM name</i> exists, are you going to over-write it?
-----------------	--

GVM0012W	The VM <i>VM name</i> is running, make sure the system is powered down, then hit OK to continue.
-----------------	---

GVM0020E	A connection with the vCenter server could not be established.
-----------------	---

Explanation:

The server might not be running.

Administrator response:

This might indicate a network problem. Ensure that the server is running and the machine is accessible. Try the action again.

GVM0021I **A connection with the vCenter server has been established.**

GVM0022E **The VMCLI inquire configuration command failed, the following messages describe the error.**

Explanation:

The Derby database might not be running.

Administrator response:

Correct the problem. Try the action again.

GVM0023I **The VMCLI inquire configuration command completed successfully.**

GVM0024E **Failed to determine which product or products are installed.**

Explanation:

See message.

Administrator response:

Correct the problem. Try the action again.

GVM0025I **Successfully determined which product or products are installed.**

GVM0026E **Multiple restore points have been selected, but they are not located in the same datacenter.**

Explanation:

Selecting restore points from different datacenters is not permitted. The restore points must all be located in the same datacenter.

Administrator response:

Select the restore points from the same datacenter or select just a single restore point.

GVM0027E **Multiple restore points have been selected, but they are not from the same backup.**

Explanation:

Selecting restore points from different backups is not permitted. The restore points must all be located in the same backup.

Administrator response:

For restores from IBM Storage Protect Snapshot, all restore points must come from the same backup. You cannot restore multiple VMs that come from of different backups.

GVM0028E **A key configuration file is missing: vmcliConfiguration.xml.**

Explanation:

The file vmcliConfiguration.xml is required for the GUI to operate, but has not been found during GUI session

startup. This is an unusual problem, it may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the file is located in the correct directory, has correct access permissions, and has valid syntax for its content. Retry accessing the GUI.

GVM0029E **Invalid mode tag in file vmcliConfiguration.xml.**

Explanation:

The xml tag mode in file vmcliConfiguration.xml is required for the GUI to operate, but is missing or has an incorrect value. This may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM0030E **Invalid enable_direct_start tag in file vmcliConfiguration.xml.**

Explanation:

The xml tag enable_direct_start in file vmcliConfiguration.xml is required for the GUI to operate, but is missing or has an incorrect value. This may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM0031E **Invalid URL tag for the specified mode tag in file vmcliConfiguration.xml.**

Explanation:

In file vmcliConfiguration.xml, the URL tag corresponding to the specified mode tag is required for the GUI to operate, but is missing or has an incorrect value. This may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the correct URL tag is specified with a valid value for the specified mode. Retry accessing the GUI.

GVM0032E **Invalid VMCLIPath tag in file vmcliConfiguration.xml.**

Explanation:

The xml tag VMCLIPath in file vmcliConfiguration.xml is required for the GUI to operate, but is missing or has an incorrect value. This may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM0033E **Invalid interruptDelay tag in file vmcliConfiguration.xml.**

Explanation:

The xml tag interruptDelay in file vmcliConfiguration.xml is required for the GUI to operate, but is missing or has an incorrect value. This may be due to an install issue or manual editing of the file.

Administrator response:

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM0099E	The VM name entered <i>VM name</i> conflicts with an existing VM. Please enter a different name.
GVM0100E	An error occurred while processing the request to the Web server. If this error persists, check the network connection with the Web server and verify that the Web server is running.Detail: <i>exception exception message</i>
GVM0101E	A request to the server took too long to complete. If this error persists, check the network connection with the Web server and verify that the Web server is running.
GVM0102E	An error occurred while processing the response from the Web server.Detail: <i>error</i>
GVM0103E	An error occurred while making the Web server request. If this error persists, check the network connection with the Web server and verify that the Web server is running.Error: <i>message</i>
GVM0104E	No matching device class found. Please return to source page and reselect.
GVM0105E	No matching proxy node found. Please return to source page and reselect.
GVM0106E	No proxy ESX hosts available.
GVM0107I	Password set successfully.
GVM0108E	Set password failed.Error: <i>message</i>

Explanation:

The password may be incorrect or the server is not running.

Administrator response:

Verify the password is correct then try the action again. Or check the network connection with the

server machine and verify that the server is running then try the action again.

GVM0109E	Get managed domain failed.Error: <i>message</i>
-----------------	--

GVM0110E	Multiple restore points have been selected, but they are not the same backup type.
-----------------	---

Explanation:

Selecting restore points of different types is not allowed. The restore points must all be located on either a IBM Storage Protect server or in the IBM Storage Protect Snapshot repository.

Administrator response:

Select the same type of restore points or select just a single restore point.

GVM0111E	Backup ID is null.
-----------------	---------------------------

Explanation:

An internal error occurred.

Administrator response:

Refresh the table and perform the action again.

GVM0112E	Task ID is null.
-----------------	-------------------------

Explanation:

An internal error occurred.

Administrator response:

Refresh the table and perform the action again.

GVM0113E	Could not open a pop-up window.
-----------------	--

Explanation:

An internal error occurred.

Administrator response:

Try the action again.

GVM0114E	Virtual machine name is null.
-----------------	--------------------------------------

Explanation:

An internal error occurred.

Administrator response:

Refresh the table and perform the action again.

GVM0115E	Datastore does not exist.
-----------------	----------------------------------

Explanation:

An internal error occurred.

Administrator response:

Refresh the table and perform the action again.

GVM0116I	No selection was made, the whole virtual machine will be attached.
-----------------	---

Explanation:

No selection was made.

Administrator response:

Continue with the action or cancel the action.

GVM0117I	Domain set successfully.
GVM0118E	Set domain failed.Error: <i>message</i>
Explanation: The server might not be running. The permissions on the file directory may be incorrect.	
Administrator response: Check the network connection with the server machine. Verify that the server is running and try the action again. Check the permissions of the directory indicated in SystemErr.log if error indicates incorrect permissions.	
GVM0119E	The schedule requires use of the following datacenters that are not in the active domain. Datacenters: <i>list</i> Action: This schedule may not be updated, instead either update the domain construct to include the datacenters, or create a new schedule without dependence on these datacenters. Detail: The schedule definition is as follows: Schedule Summary: <i>summary</i>
GVM0120E	The schedule requires use of the following datacenters that are not known to the system. Datacenters: <i>list</i> Action: This schedule may not be updated, instead create a new schedule without dependence on these datacenters. Detail: The schedule definition is as follows: Schedule Summary: <i>summary</i>
GVM0121E	The schedule requires use of the following hosts that are not known to the system. Hosts: <i>list</i> Action: This schedule may not be updated, instead create a new schedule without dependence on these hosts. Detail: The schedule definition is as follows: Schedule Summary: <i>summary</i>
GVM0122E	The schedule requires use of the following datastores that are not known to the system. Datastores: <i>list</i> Action: This schedule may not be updated, instead create a new schedule without dependence on these datastores. Detail: The schedule definition is as follows: Schedule Summary: <i>summary</i>
GVM0123E	The schedule requires use of the following virtual machines that are not known to the system. Virtual Machines: <i>list</i> Action: This

schedule may not be updated, instead create a new schedule without dependence on these virtual machines. Detail: The schedule definition is as follows: Schedule Summary: *summary*

GVM0124I	Password set successfully. Warning: <i>message</i>
-----------------	---

Explanation:
The password was set successfully with a warning.

Administrator response:
Follow the action described in the warning message.

GVM0125E	An error occurred while making the Web server request. If this error persists, check the network connection with the Web server and verify that the Web server is running. Error: <i>error</i>
-----------------	---

GVM1100E	The following command requires confirmation from the server: <i>Command</i>
-----------------	--

Explanation:
A command was issued, and a reply was expected. Some commands require a confirmation, which you cannot issue through the IBM Storage Protect GUI.

Administrator response:
Issue the command from the command line.

GVM1101E	The following command is unknown to the server: <i>Command</i>
-----------------	---

Explanation:
An unknown command was issued to the server. The command might not be valid on the server version and platform or the command syntax might be incorrect.

Administrator response:
Verify that the command is valid for the server version and platform, and verify that the command syntax is correct.

GVM1102E	The syntax of the following command is incorrect: <i>Command</i>.
-----------------	--

Explanation:
See message.

Administrator response:
Correct the syntax and issue the command from the command line. The activity log of the IBM Storage Protect Server shows all the commands issued before and after this command.

GVM1103E	An internal server error occurred.
-----------------	---

Explanation:
See message.

Administrator response:

Try the command again. If this does not work, contact customer support. You might be asked to provide tracing information and information about the actions performed before the failure occurred.

GVM1104E **The server ran out of memory while processing the request. Close any unnecessary processes on the IBM Storage Protect server and try the operation again.**

Explanation:
See message.

Administrator response:
Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1105E **The database recovery log is full.**

Explanation:
See message.

Administrator response:
Before trying the action again, extend the recovery log or back up the IBM Storage Protect server database. Contact the administrator of the IBM Storage Protect server.

GVM1106E **The server database is full.**

Explanation:
See message.

Administrator response:
Before trying the action again, extend the server database. Contact the administrator of the IBM Storage Protect server.

GVM1107E **The server is out of storage space.**

Explanation:
See message.

Administrator response:
Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1108E **You are not authorized to perform this action. An administrator with system authority can change your authority level to allow you to perform this action.**

GVM1109E **The object that you are attempting to access does not exist on the server.**

GVM1110E **The object that you are attempting to access is currently in use by another session or process. Retry the action at a later time.**

GVM1111E **The object that you are attempting to remove is referenced by another object defined to the server.**

Remove the other object before removing this one.

GVM1112E **The object that you are attempting to access or remove is not available.**

Explanation:
See message.

Administrator response:
Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1113E **The server encountered an I/O error while processing the request. For more information, see the operating system event or error log.**

GVM1114E **The action failed because the transaction could not be committed.**

Explanation:
See message.

Administrator response:
Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1115E **The action failed because of a resource lock conflict.**

Explanation:
See message.

Administrator response:
Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1116E **The action failed because of a mode conflict.**

Explanation:
See message.

Administrator response:
Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1117E **The action failed because the server could not start a new thread.**

Explanation:
See message.

Administrator response:
Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1118E	The server is not licensed to perform this action. If a license was purchased, use the command line to register the license.
-----------------	---

GVM1119E	The specified destination is not valid.
-----------------	--

Explanation:
See message.

Administrator response:
Enter a different destination or update the configuration with a valid destination, and try the action again.

GVM1120E	The specified input file cannot be opened. Verify the file name and directory permissions, then try the action again.
-----------------	--

GVM1121E	The specified output file cannot be opened. Verify the file name and directory permissions, then try the action again.
-----------------	---

GVM1122E	An error occurred while writing to the specified output file.
-----------------	--

Explanation:
See message.

Administrator response:
Check the file system to ensure that there is enough space. Check the operating system event or error log for more information.

GVM1123E	The specified administrator is not defined to this server.
-----------------	---

Explanation:
See message.

Administrator response:
Ensure that the administrator name was entered correctly. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1124E	The SQL statement could not be processed.
-----------------	--

Explanation:
An exception occurred while processing the SQL statement. Possible exceptions include divide-by-zero, math overflow, temporary table storage space unavailable, and data-type errors.

Administrator response:
Correct the SQL query and try again.

GVM1125E	This operation is not allowed with this object.
-----------------	--

Explanation:
See message.

Administrator response:
Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1126E	The table was not found in the server database.
-----------------	--

Explanation:
See message.

Administrator response:
Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1127E	The specified file space name is not compatible with the filespace type.
-----------------	---

Explanation:
Unicode file space names are incompatible with non-unicode names.

Administrator response:
Enter a file space name of the correct type and try the action again.

GVM1128E	The specified TCP/IP address is not valid. Verify the TCP/IP address and try the action again.
-----------------	---

GVM1129E	No objects were found that match the search conditions.
-----------------	--

GVM1130E	Your administrative ID on this server is locked. An administrator with system authority can unlock your ID.
-----------------	--

GVM1131E	The connection to the server was lost while performing the action.
-----------------	---

Explanation:
See message.

Administrator response:
This might indicate a network problem. Ensure that the server is running and the machine is accessible. Retry the action.

GVM1132E	Your ID or password is not valid for this server.
-----------------	--

Explanation:
See message.

Administrator response:
Launch the Configuration Editor from the Configuration Tab and enter a valid ID or password for your IBM Storage Protect Server.

GVM1133E	Your password expired on this server.
-----------------	--

Explanation:
Your IBM Storage Protect password has expired.

Administrator response:

Reset your password on the IBM Storage Protect Server or contact your IBM Storage Protect Server administrator to reset it.

GVM1134E	The server cannot accept new sessions. If sessions are disabled for this server, issue the ENABLE SESSIONS command from the command line.
GVM1135E	A communications failure occurred while processing the request. Retry the action at a later time.
GVM1136E	The administrative API encountered an internal error while processing the request.
GVM1137E	The administrative API cannot process the command document sent from the server.

Explanation:

The XML command document could not be parsed. Either the file could not be read, or the file is corrupted.

Administrator response:

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1138E	The following command contains one or more invalid parameters: <i>command</i>.
-----------------	---

Explanation:

The IBM Storage Protect GUI tried to run a command, but the call to the API contained one or more invalid parameters.

Administrator response:

Check the parameters in the command. If you entered text in a field, you might find the error in the parameters and correct it. Viewing the activity log might help to determine the cause of the problem. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1139E	The administrative API encountered invalid parameters while processing the request.
-----------------	--

Explanation:

A command was run through the administrative API, but one of the parameters to an API method was invalid.

Administrator response:

This is typically an internal error, but it can be caused by unusual parameters. For example, characters such as: < > & can cause the problem. Check the parameters in the command. If you entered text in a field, you might find the error in the parameters and correct it.

GVM1140E	The administrator's authority level on this server cannot be determined.
-----------------	---

Explanation:

See message.

Administrator response:

Use a different administrator ID. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM1141E	An object with the name that you specified already exists on the server. Enter a different name.
-----------------	---

GVM1142E	The version of the server is not supported by the IBM Storage Protect GUI.
-----------------	---

GVM1143E	An internal error has occurred.
-----------------	--

Explanation:

The operation failed after encountering an internal error.

Administrator response:

Retry the operation. If this does not work, contact customer support. You might be asked to provide tracing information and information about the actions performed before the failure occurred.

GVM1144E	The operation failed, please go to the log for more details.
-----------------	---

GVM1145E	Wrong format of the end date and time. Please enter the end date and time format as yyyyMMddHHmmss.
-----------------	--

GVM1146E	Sorry, the description of the backup task was not created in a file. Please try again.
-----------------	---

Explanation:

On the general page of the backup wizard, you can describe your backup task in general.

GVM1147E	The ESXHOST name you entered is too long. Please change to a shorter one.
-----------------	--

GVM1148E	Wrong Backup ID. Please try again.
-----------------	---

GVM1150E	An error occurred when processing the backup object file. Please try again later.
-----------------	--

Explanation:

When you click submit in the backup wizard, the object list will be stored in a file. When processing this file, an error occurred.

GVM1151E	No backup object is selected. You must choose a source node to backup.
Explanation: To initiated a backup task, you have to choose an object on the source page of the backup wizard.	
GVM1152E	Wrong format of the start date and time. Please enter the start date and time format as yyyyMMddHHmmss.
GVM1153I	Backup task <i>Task Name</i> started, would you like to monitor this task now?
GVM1154I	Delete backup task completed successfully.
GVM1155E	Delete backup task failed, please check log for more detail.
GVM1156I	Restore Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
GVM1157E	<i>Error Or Warning</i>
GVM1158I	Mounted backup Item could not be restored.
GVM1159I	Result of attach is <i>status</i> (Task ID: <i>Task ID</i>), refer to events list to get the details.
GVM1160I	Result of detach is <i>status</i> (Task ID: <i>Task ID</i>), refer to events list to get the details.
GVM1161I	Command successfully submitted to the IBM Storage Protect server. Detail: <i>Server Messages</i>
GVM1162E	The command submitted to the IBM Storage Protect server failed. Error: <i>Error Code Error Messages</i>

Explanation:

The cause of the problem is identified in the message text.

Administrator response:

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM1163E	No IBM Storage Protect server connection, please configure the IBM Storage Protect server in the configuration panel.
GVM1164E	The selected items can only be under ONE datacenter.

GVM1165E	Authentication failed. Could not connect to vCenter. Make sure you log in using the VMware vSphere client and have a valid session.
GVM1166E	Authentication failed. Please log in using the VMware vSphere client.
GVM1167E	The virtual machine <i>VM name</i> exists. Delete the virtual machine first before restoring it.
GVM1168E	The target virtual machine <i>VM name</i> is running. Close the virtual machine before restoring virtual disks to it.
GVM1169E	Some of selected virtual disks exist in target virtual machine. Remove those virtual disks from target virtual machine before restoring to it.
GVM1170E	A VMCLI command failed. Error: <i>Error Messages</i>

Explanation:

The cause of the problem is identified in the message text.

Administrator response:

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM1171E	A request submitted to the VMware vCenter server failed. Error: <i>Error Messages</i>
-----------------	--

Explanation:

The cause of the problem is identified in the message text.

Administrator response:

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM1172E	A command submitted to the IBM Storage Protect server failed. Error: <i>Error Messages</i>
-----------------	---

Explanation:

The cause of the problem is identified in the message text.

Administrator response:

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM1173E	Cannot find the file with format 'summary.date.log' in the path: <i>path</i>
GVM1174E	Cannot find the IBM Storage Protect Snapshot installation path using the VMCLI inquire_config command.
GVM1175E	A VMCLI command to get version failed.
GVM1176I	Backup task <i>Task ID</i> started, would you like to monitor this task now?
GVM1177E	The IBM Storage Protect Web Server could not be contacted.

Explanation:

The IBM Storage Protect GUI has attempted to contact its Web Server. The operation was not successful.

Administrator response:

Perform one or more of the following steps to try and determine the problem: Verify that the IBM Storage Protect Web Server is running. Verify that the Web Server machine is running. Verify that the Web Server machine is accessible over the network. Close the IBM Storage Protect GUI. Start the GUI again when the problem is resolved.

GVM1178I	Command successfully submitted to the server.
GVM1179E	No host is found in datacenter <i>datacenter name</i>. Select another datacenter to restore.
GVM1180W	The schedule does not contain all the required parameters. It cannot be displayed in the properties notebook.

Explanation:

This schedule may have been created or modified outside of IBM Storage Protect GUI.

Administrator response:

This schedule must be modified outside the IBM Storage Protect GUI.

GVM1181W	One or more VMs exist. Do you want to continue the restore operation and overwrite the existing VMs?
GVM1182E	The Administrator Id provided does not have sufficient privileges.

Explanation:

The operation you are attempting requires a IBM Storage Protect Server Administrator Id to have at least Unrestricted Policy privilege.

Administrator response:

Contact your IBM Storage Protect Server Administrator to grant you Unrestricted Policy privilege for your Administrative Id. Or, use an alternate Id with such privilege and try again.

GVM1183E	The nodename <i>node name</i> is already in use. Please choose another nodename.
-----------------	---

Explanation:

The node name chosen already exists on the server. Choose another name.

Administrator response:

Pick another node name to use. If you want to re-use this node, then unselect the 'Register Node' checkbox.

GVM1184E	The node name <i>node name</i> is not defined on server. Make sure the node name you entered exists on the server.
-----------------	---

Explanation:

The node name entered does not exist on the server. Since you did not select 'Register Node' checkbox, the node name you enter must have been previously defined and exist on the server.

Administrator response:

Check the node name you are supposed to use and enter it again. If you want to register this node, then select the 'Register Node' checkbox.

GVM1185E	The passwords in the entry field and the verify field do not match. Please try again.
-----------------	--

Explanation:

The new passwords entered do not match.

Administrator response:

Clear the fields and enter the same password in both password fields.

GVM1186W	Please select one or more Datacenters to be managed.
-----------------	---

Explanation:

At least one Datacenter must be selected.

Administrator response:

Add one or more Datacenter(s) into the Managed Datacenters list.

GVM1187W	One or more nodes do not have their password set. Make sure all nodes have their password set.
-----------------	---

Explanation:

If a node has 'Register Node' checkbox set, then that node's password must be set.

Administrator response:

Assign a password for nodes that are to be registered.

GVM1188I	No datacenter node was found mapped to <i>datacenter name</i> . Select a datacenter node from the list to associate with <i>datacenter name</i> . Leave the selection empty to have the Configuration Wizard create a new datacenter node for it.
-----------------	---

GVM1189I	Are you sure you want to proceed without entering a IBM Storage Protect Administrative ID? Without IBM Storage Protect Administrative access, the Wizard will not validate node names or register nodes. Instead, a macro file will be generated at the end of this Wizard for you to give to your IBM Storage Protect Administrator to execute.
-----------------	--

GVM1190I	This task was skipped because it was not necessary or a prerequisite task failed.
-----------------	---

GVM1191E	There was an error writing to script file: <i>file path</i> .
-----------------	---

Explanation:

An error was encountered when trying to write to file at the path indicated.

Administrator response:

Try the operation again.

GVM1192I	Managed datacenters have changed. Please go to the data mover page to verify or change your current mappings.
-----------------	---

GVM1193I	No datacenter nodes were found for the vCenter node <i>vCenter node</i> and VMCLI node <i>VMCLI node</i> configuration. The Wizard will generate a default set of datacenter nodes for you.
-----------------	---

GVM1194E	The password entered is not acceptable. Choose another password.
-----------------	--

Explanation:

IBM Storage Protect Server could not accept the password chosen. It could be because the password did not meet certain password rules.

Administrator response:

Try with another password.

GVM1195W	Unchecking this checkbox means you are supplying a node name that is already defined on the IBM Storage Protect Server AND that
-----------------	---

it is meant to be used for your configuration. Since this Wizard is proceeding without Administrative access, it cannot verify if the node exists or not. You should only proceed if you understand what you are doing.

Explanation:

Since you are using the Configuration Wizard without a IBM Storage Protect Administrative ID, you should be very careful. The macro script file generated at the end of running the Configuration Wizard could contain errors because values are not validated.

Administrator response:

We strongly recommend you use the Configuration Wizard with a proper IBM Storage Protect Administrative ID.

GVM1196W	The IBM Storage Protect node <i>node</i> has already been identified. If you want a different name other than the default name, edit this field again. If you want to use the same data mover for multiple Datacenters, please use Configuration Settings to do this.
-----------------	---

Explanation:

The node is already being used in this configuration.

Administrator response:

Try using another node name.

GVM1197W	The IBM Storage Protect node <i>node</i> has invalid characters or exceeds 64 characters. Choose a different name and edit this field again.
-----------------	--

Explanation:

The node name is invalid or longer than 64 characters.

Administrator response:

Try using another node name.

GVM1198E	The password entered is not acceptable on this Server because it contains invalid characters. The valid characters are: <i>validCharsString</i>
-----------------	---

Explanation:

IBM Storage Protect Server could not accept the password chosen because of invalid characters in the password.

Administrator response:

Try with another password that only contain valid characters.

GVM1199E	The password entered is not acceptable on this Server because
-----------------	---

of the reason below. Choose another password. Error: *message*

Explanation:

IBM Storage Protect Server could not accept the password chosen. The reason why this password is not valid is given in the message.

Administrator response:

Try with another password that meets the rules.

GVM1200E **Filter has changed, select Apply filter before continuing.**

Explanation:

Filter pattern must be applied after it is changed.

Administrator response:

Click the Apply filter button.

GVM1201E **Select at least one item from a datacenter to continue.**

Explanation:

A host, host cluster, or VM must be selected to do a backup.

Administrator response:

Select an item under a datacenter.

GVM1202E **Your selections exceed the 512 character limit allowed for backups, change your selection.**

Explanation:

The number of characters required to list the selected items exceeds the limit of 512 characters. Also, if hosts have been partially selected, characters are needed to list the VMs that are excluded from the backup.

Administrator response:

Create multiple backup tasks, with less selected items per task.

GVM1203I **Changing the newly added virtual machines checkbox clears all selections of host clusters, hosts, and virtual machines. Press OK to proceed, or Cancel to leave unchanged.**

Explanation:

The state of the newly added virtual machines checkbox significantly impacts what is allowed to be selected on the source panel, so selections are cleared when the state changes.

Administrator response:

Select OK to proceed, or select Cancel to retain all selections.

GVM1204E **Datacenter node *datacenter node name* does not have a IBM Storage**

Protect node mapped in the vmcli configuration file.

Explanation:

The datacenter node must have a corresponding IBM Storage Protect node listed in the configuration file named vmcliprofile.

Administrator response:

Correct the problem by going to the Configuration tab in the GUI and selecting Edit Configuration to update the mapping for the datacenter. Also resolve any other configuration errors that are reported on the Configuration tab.

GVM1205E **IBM Storage Protect datacenter node *datacenter node name* maps to vCenter datacenter name *datacenter name* in the vmcli configuration file, but *datacenter name* does not exist in the vCenter.**

Explanation:

The vCenter datacenter name maps to a datacenter node in the vmcli configuration file named vmcliprofile, but the data enter name does not exist in the vCenter.

Administrator response:

Correct the problem by going to the Configuration tab in the GUI and selecting Edit Configuration to update the mapping for the datacenter. Also resolve any other configuration errors that are reported on the Configuration tab.

GVM1206E **You have selected items from multiple datacenters: *datacenter list*. This is not allowed, all selections must be from one datacenter.**

Explanation:

A backup task only supports items from one datacenter. If this is an existing task, changes in the vCenter configuration after task creation may have caused this problem.

Administrator response:

Check and correct the selections to make sure all selections are under the same datacenter.

GVM1207E **The selected items *item list* are not found under datacenter *datacenter name* in the vCenter, please review and de-select them.**

Explanation:

Items originally selected are no longer found under the datacenter associated with the backup task. This may be caused by changes in the vCenter configuration.

Administrator response:

Review if the items are now located under a different datacenter. De-select the not found items, and make

new selections under the other datacenter or create a new backup task for these items.

GVM1208I **The datacenter in the source page has changed, please reselect the data mover node in the Destination page.**

Explanation:

When item selection is changed to a different datacenter, the valid data movers may change. You are required to select the data mover again on the Destination page.

Administrator response:

Reselect the data mover node in the Destination Page.

GVM1209I **Are you sure you want to use node *data mover node* as a data mover for datacenter *datacenter*?**

GVM1210I **Are you sure you want to use node *node name* that is already registered on the IBM Storage Protect server as a data mover for datacenter *datacenter*? If so, we will mark the node as such and you will be unable to make any further changes to the node.**

GVM1211E **The password entered is not acceptable on this Server because it is too short. Passwords must have a least *minPasswordLength* characters.**

Explanation:

IBM Storage Protect Server could not accept the password chosen because it is too short.

Administrator response:

Try with another password that is longer than the required minimum length.

GVM1212E ***Component* is downlevel, so its use is disabled in the GUI. You will only be able to use the GUI for *component*.**

GVM1213E **Mismatching IBM Storage Protect Server entries in the current settings is detected. IBM Storage Protect Server definition used by the GUI: *server1* IBM Storage Protect Server where backups are stored: *server2* Click Reset Server definition to clear the IBM Storage Protect definition and enter new credentials. Or click on Reconfigure Environment to launch the Configuration Wizard to reconfigure your IBM Storage Protect environment.**

Explanation:

IBM Storage Protect detected mismatching IBM Storage Protect Server entries between the vmcliprofile and the current GUI's IBM Storage Protect Server connection.

Administrator response:

Pick one of the two actions available. You may either reset the IBM Storage Protect Server definition/credentials OR use the Configuration Wizard to set up a new environment.

GVM1214E **The SSL Connection could not be made. The IBM Storage Protect SSL certificate is missing. Check for valid IBM Storage Protect certificate in the TSM-ve-truststore.jks**

Explanation:

IBM Storage Protect Server did not accept the SSL connection. SSL keystore is not in the default location or does not contain a IBM Storage Protect certificate.

Administrator response:

Check the TSM-ve-truststore.jks for a valid certificate, ensure TSM-ve-truststore.jks is in the correct default location.

GVM1215E **The password entered is not acceptable on this Server because it is too long. Passwords cannot have more than *maxPasswordLength* characters.**

Explanation:

IBM Storage Protect Server could not accept the password chosen because it is too long.

Administrator response:

Try with another password that is shorter than the allowed maximum length.

GVM1216E **The SSL Connection could not be made. The IBM Storage Protect SSL certificate has expired.**

Explanation:

IBM Storage Protect Server did not accept the SSL connection. The TSM-ve-truststore.jks has an expired IBM Storage Protect SSL certificate.

Administrator response:

Obtain a new valid IBM Storage Protect SSL certificate from the IBM Storage Protect server and place it in the TSM-ve-truststore.jks.

GVM1217E **The non-SSL connection could not be made. This IBM Storage Protect Admin ID requires an IBM Storage Protect SSL connection.**

Explanation:

IBM Storage Protect Server did not accept the non-SSL connection. The IBM Storage Protect Server requires SSL be used with this Admin ID.

Administrator response:

Use SSL with this Admin ID. Ensure that the TSM-ve-truststore.jks with a valid IBM Storage Protect server SSL certificate is installed in the default location.

GVM1218E Your selections have caused the backup task definition to require *count* characters, which exceeds the 512 character limit. This can be caused by a long virtual machine exclude list, which is the list of all VMs under host(s) that were not selected. Either select more VMs under selected hosts or de-select the newly added virtual machines checkbox.

Explanation:

When the newly added virtual machines checkbox is selected, the resulting backup task must list all unselected VMs for hosts that are partially selected. The backup task definition has a 512 character limit, and the combination of selected items and excluded VMs exceeds this limit.

Administrator response:

De-select the newly added virtual machines checkbox or create multiple backup tasks with less selected items per task.

GVM1219E Your selection of virtual machines has caused the backup task definition to require *count* characters, which exceeds the 512 character limit. Either create multiple backup tasks with less virtual machines per task, or select the newly added virtual machines checkbox and choose entire hosts with no more than a few unselected VMs.

Explanation:

The backup task definition has a 512 character limit, and the total number of characters for the selected items exceeds this limit.

Administrator response:

Create multiple backup tasks with less selected virtual machines per task, or select the newly added virtual machines checkbox and then select hosts instead of individual virtual machines (you can de-select a small number of virtual machines per host if desired.)

GVM1220E There is no data mover node proxy relationship for datacenter node *datacenter node name*. Review the data mover relationships on

the Configuration tab or the IBM Storage Protect server.

GVM1221E There is no datacenter node defined for datacenter *datacenter name*. Review the node configuration on the Configuration tab.

GVM1222I Node *name name* is currently locked. The Configuration Wizard will attempt to unlock this node if you choose to continue.

GVM1223E A connection with the IBM Storage Protect server (*Address:Port*) could not be established. Please verify the server address and *Server or Admin port* are correct.

Explanation:

The server might not be running or specified admin port or server port may be incorrect.

Administrator response:

Check the network connection with the IBM Storage Protect server machine. Verify that the server is running and try to log in again. Also verify server port and admin port information is correct.

GVM1224E The vCenter user name or password is not valid. Please try again.

Explanation:

The vCenter user name or password is not valid.

Administrator response:

Enter the user name or password again.

GVM1225E Permission to perform this operation was denied. Please try with other user name.

Explanation:

The vCenter user name is not valid.

Administrator response:

Enter another user name.

GVM1250I A IBM Storage Protect Administrative ID and password is currently not set. The absence of this information limits the actions that you can take in the GUI. Click OK to be taken to the configuration settings panel and enter an ID and password. Click Cancel to continue without using an ID and password.

GVM1251W You have chosen an Administrative ID that has less authority than the current ID. Are you sure you want change this

ID? Current IBM Storage Protect Authority Level: *Current Level* New IBM Storage Protect Authority Level: *New Level* Current Role: *Current Role* New Role: *New Role* Click OK to accept these changes, or Cancel to exit without change.

GVM1252I	Here are the current and new roles for IBM Storage Protect Admin IDs. Review and confirm these changes. Current IBM Storage Protect Authority Level: <i>Current Level</i> New IBM Storage Protect Authority Level: <i>New Level</i> Current Role: <i>Current Role</i> New Role: <i>New Role</i> Click OK to accept these changes, or Cancel to exit without change.
GVM1253I	ID has been changed without save. Previous ID will be loaded.
GVM1254I	Your current UI role does not allow you to unlock or reset the VMCLI node. In order to make changes, go to the Server Credentials page and enter a IBM Storage Protect Admin ID and password that has the necessary privileges for making VMCLI node updates. Select OK to save these credentials, then re-open the Configuration Settings notebook and you can make VMCLI node updates.
GVM1255I	Your current UI role does not allow you visit other panels. Select OK to save these credentials, then re-open the Configuration Settings notebook and you can make other updates.
GVM1256I	There are non-English characters contained in one or more datacenters. The domain will be adjusted accordingly.
GVM1257E	Datacenter <i>DataCenter Name</i> cannot be added to the domain because it contains non-English characters.

Explanation:

Datacenters that contain non-English characters are not currently supported. Therefore, they cannot be added to the domain.

Administrator response:

Datacenter will not be added to the domain.

GVM1258W Node *Node Name* already exists on the server. Attempt to rename node to *New Node Name*?

Explanation:

Node name is already registered on the IBM Storage Protect server.

Administrator response:

Click Yes to attempt to rename node. Click No to make other changes. Example: unclick register node, rename node manually.

GVM1259W The following virtual machines for host *Host Name* have unsupported characters in their name: *Invalid Virtual Machine Names*. Therefore, these virtual machines are not backed up, regardless of your selections. You must rename these virtual machines to back them up.

Explanation:

The following characters are not supported in virtual machine names: ' ; * ? , < > / \ |

Administrator response:

Rename the identified virtual machines to remove unsupported characters from their name.

GVM1260E The following host clusters have unsupported characters in their name: *Invalid Host Clusters*. These host clusters cannot be selected for backup because they contain unsupported characters. Rename these host clusters or remove them from selection.

Explanation:

The following characters are not supported in host cluster names: ' ; * ? , < > / \ |

Administrator response:

Rename the identified host clusters to remove unsupported characters from their name. Or, remove them from your backup selection.

GVM1261E Your selections created an empty virtual machine list for backup. This issue might occur because all the selected virtual machines contain unsupported characters in their names. Make sure that you selected virtual machines that do not contain unsupported characters in their names.

Explanation:

The following characters are not supported in virtual machine names: ' ; * ? , < > / \ | . Virtual machine names that contain these characters are automatically

removed from the backup task definition. This removal can cause an empty task definition.

Administrator response:

Rename the identified virtual machines to remove unsupported characters from their name. Or, select different virtual machines to back up.

GVM1262E	The filter pattern cannot be applied because it contains unsupported characters. Change the pattern to remove the unsupported characters, then apply the filter again.
-----------------	---

Explanation:

The following characters are not supported in filter pattern: ' : ; < > / \ |

Administrator response:

Change the filter pattern to remove unsupported characters, then apply the filter again.

GVM1263E	A temporary datastore is not available to perform this operation. This temporary datastore is required in addition to the restore destination datastore.
-----------------	---

Explanation:

A datastore is required for use as a temporary restore destination for this operation. This temporary datastore must be from the same ESX host as the datastore that is used for the actual restore destination. However, the temporary datastore cannot be the same datastore that is used for the actual restore destination.

Administrator response:

Add a datastore to the destination ESX host. Then, select this datastore as the temporary restore destination.

GVM1264E	There was an error creating opt file: <i>file name</i>.
-----------------	--

Explanation:

An error was encountered when trying to write to file.

Administrator response:

Try the operation again.

GVM1265E	Creating <i>service</i> has failed. No services were created for data mover node <i>node name</i>.
-----------------	---

Explanation:

An error was encountered when trying to create IBM Storage Protect service for data mover node specified.

Administrator response:

Check environment and ensure user has proper rights before trying operation again.

GVM1266E	Creating firewall for <i>service</i> has failed. Please manually add firewall rules for services installed.
-----------------	--

Explanation:

An error has occurred when attempting to add firewall rule for specified executable.

Administrator response:

Check environment and ensure user has proper rights before trying operation again or manually add rule to firewall for IBM Storage Protect client acceptor, IBM Storage Protect Agent and IBM Storage Protect Scheduler.

GVM1267W	Local services were setup successfully but were unable to verify firewall access for these executable files: <i>agentExe cadExe schedExe</i> If any problems are experienced related to local services, verify that firewall access is available for these executable files.
-----------------	---

Explanation:

Microsoft firewall may be disabled or another firewall may be in place.

Administrator response:

Check environment and add rules manually if needed for the IBM Storage Protect client acceptor, IBM Storage Protect Agent, and IBM Storage Protect Scheduler.

GVM1269E	Reason Code <i>reason</i> This error was reported by the IBM Storage Protect data mover. No further description is available. For more information, review the error log <i>errorLog</i> on the data mover host machine <i>hostname</i> at address '<i>address</i>'.
-----------------	---

Explanation:

The data mover encountered an error with the reported reason code.

Administrator response:

Log into the host machine specified and view the error log for more information.

GVM1270W	Warning: If this task is canceled, all created data on the virtual machines that are not completely restored is lost and the virtual machines are removed from the ESX host. Are you sure that you want to cancel this task?
-----------------	---

Explanation:

A cancel task command is submitted. Refresh to see the cancel progress.

Administrator response:

Cancel the selected task or allow the task to continue processing.

GVM1271W	Scan schedule <i>schedule name</i> was successfully defined on the server and associated with node <i>node name</i>, however no services were created to run the schedule. Detail: <i>error</i>
-----------------	--

Explanation

An error was encountered in one of the steps below when trying to create IBM Storage Protect services for the VMCLI node.

1. Create the option file for the VMCLI node.
2. Set the password for the VMCLI node to a temporary password for the next step.
3. Run the IBM Storage Protect Client Service Configuration Utility to create the services.
4. Run the IBM Storage Protect Client Service Configuration Utility to start the client acceptor service.
5. Reset the VMCLI node password.

Administrator response:

Delete the schedule and create the schedule again to automatically configure the services or manually configure the services. Check environment and ensure user has proper rights before trying operation again.

GVM1272W	Scan schedule <i>schedule name</i> was successfully defined on the server and associated with node <i>node name</i>. IBM Storage Protect services were created to run the schedule. However, resetting the VMCLI node password failed. Detail: <i>error</i>
-----------------	--

Explanation:

An error was encountered while trying to reset the VMCLI node password.

Administrator response:

Use the Configuration Settings to reset the VMCLI node password.

GVM1273W	A dismount operation removes the iSCSI disks but does not remove the VM or its data. Before proceeding with dismount, make sure the following conditions exist: -The mounted iSCSI disk is recovered. -Storage vMotion completed migrating the VM to a
-----------------	---

local datastore. If the recovery operation failed and you want to delete the VM, its data, and dismount any iSCSI targets, click Dismount and Delete. Dismount and Delete is a destructive action and deletes the VM and its data, regardless of the success or failure of the instant restore operation. Based on this information, do you want dismount the VMs that are selected for instant restore?

Explanation:

A dismount operation removes the iSCSI disks but does not remove the VM or its data. Before proceeding with dismount, make sure the following conditions exist: The mounted iSCSI disk is recovered, Storage vMotion completed migrating the VM to a local datastore. If the recovery operation failed and you want to delete the VM, its data, and dismount any iSCSI targets, click Dismount and Delete. Dismount and Delete is a destructive action and deletes the VM and its data, regardless of the success or failure of the instant restore operation.

Administrator response:

Click 'Dismount' to dismount the virtual machines that are selected for the instant restore operation. Click 'Dismount and Delete' to dismount the virtual machines that are selected for the instant restore operation, remove them from the ESX host, and verify that Storage vMotion is not running.

GVM1274W	During a dismount operation, all created data on the virtual machines is lost and the virtual machines are removed from the ESX host. Dismount the selected Instant Access virtual machines?
-----------------	---

Explanation:

All created data on the virtual machines is lost and the virtual machines are removed from the ESX host.

Administrator response:

Click 'Dismount' to dismount (cleanup) the instant access virtual machines.

GVM1275E	Selecting multiple virtual machines with different restore types is not allowed.
-----------------	---

Explanation:

Restoring multiple virtual machines with different restore types is not supported.

Administrator response:

Select virtual machines that have the same restore type.

GVM1276I	Cleanup Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
-----------------	--

GVM1277W	Are you sure that you want to cancel this task?
-----------------	--

Explanation:

A cancel task command is submitted. Refresh to see the cancel progress.

Administrator response:

Cancel the selected task or allow the task to continue processing.

GVM1278I	Your current UI role does not allow you to view backup property notebook.
-----------------	--

GVM1279I	Your current UI role does not allow you to edit nodes. In order to make changes, open the Configuration Settings notebook and go to the Server Credentials page and enter a IBM Storage Protect Admin ID and password that has the necessary privileges for making node updates.
-----------------	---

GVM1280E	Reason Code <i>reason</i> This error was reported by the IBM Storage Protect data mover. No further description is available. For more information, review the error log 'dsmerror.log' on the data mover host machine.
-----------------	--

Explanation:

The data mover encountered an error with the reported reason code.

Administrator response:

Log into the host machine where data mover resides and view the error log for more information.

GVM1281W	Login information for vCenter needed.
-----------------	--

Explanation:

In order to install new local dm services, vCenter credentials are needed.

Administrator response:

Enter vCenter credentials in order to continue.

GVM1282E	You do not have the privileges required to access the GUI.
-----------------	---

Explanation:

In order to access GUI content, the user must have the necessary vSphere privileges.

Administrator response:

Add the required privileges for the user.

GVM1283E	You do not have the permissions required to access the GUI.
-----------------	--

Explanation:

In order to access GUI content, the user must have the necessary vSphere permissions.

Administrator response:

Add the required permissions for the user.

GVM1284I	A new data center (<i>name</i>) was detected. Go to the Data Mover Nodes page to add a data center node for it.
-----------------	--

GVM1285W	The following shares and mounts will be removed and that data in there will be no longer accessible to the end user. Dismount the selected shares and mounts? <i>mounts</i>
-----------------	---

Explanation:

The selected shares and mounts will be removed.

Administrator response:

Click 'Dismount' to dismount (cleanup) the mounts and shares.

GVM1286I	Dismount Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
-----------------	---

GVM1287W	An error was encountered during the delete operation for option file: <i>file name</i>.
-----------------	--

Explanation:

An error was encountered during the delete operation. For example, this error might be caused by insufficient user permissions or the file no longer exists.

Administrator response:

Make sure the option file was deleted. If it still exists, delete this file manually.

GVM1288W	The remove operation for IBM Storage Protect service: <i>service</i> failed.
-----------------	---

Explanation:

An error prevented the IBM Storage Protect service from being removed.

Administrator response:

Check the environment and ensure that the user has sufficient rights to run this operation. Then, try the operation again.

GVM1289E	Fail to start iSCSI for mount proxy node <i>node name</i>.
-----------------	---

Explanation:

An error was encountered when trying to start iSCSI service for mount proxy node specified.

Administrator response:

Start the iSCSI service manually.

GVM1290E **The connection to the IBM Storage Protect server was not successful because either the server credentials are invalid or an SSL certificate is required but could not be obtained.**

Explanation:

A correct server user ID and password and an SSL certificate for the IBM Storage Protect server are required to connect to the server.

Administrator response:

Go to the Configuration > Tasks > Edit IBM Storage Protect Configuration > Server Credentials notebook page. Confirm that the login credentials are correct, that the correct port number is entered for the IBM Storage Protect admin port, and that the Use SSL... check box is selected. The server certificate must be retrieved and a truststore created using the procedure that is documented in the [Learn more...](#) link.

GVM1291W **The connection to the IBM Storage Protect server was not successful because a security certificate is required.**

Explanation:

Secure connections to the IBM Storage Protect server require an SSL certificate to create the connection. No certificate was found for the selected IBM Storage Protect server.

Administrator response:

If this message was not presented as part of using the configuration wizard or notebook, then proceed to the configuration notebook by choosing Configuration > Tasks > Edit IBM Storage Protect Configuration > Server Credentials.

GVM1300E **An error occurred when connecting to the vCenter server.**

Explanation:

An error occurred when connecting to the vCenter server. No further information is available.

Administrator response:

Retry the operation. If the problem persists, contact your IBM Storage Protect administrator.

GVM1301E **The vCenter user does not have the privileges that are required to perform registration operations.**

Explanation:

To perform registration operations, the vCenter user must have the Register extension and Update extension privileges, and the vCenter user is missing at least one of these privileges.

Administrator response:

Enter a vCenter user that has both required privileges, and retry the operation.

GVM1302E **An error occurred when attempting to register the plug-in to the vCenter server.**

Explanation:

An error occurred during plug-in registration. No further information is available.

Administrator response:

Retry the operation. If the problem persists, contact your IBM Storage Protect administrator.

GVM1501E **The following vcloud resources(vApp, organization, organization VDC) are invalid for selection because they have unsupported characters in their name: *reslist***

Explanation:

In order to create backup tasks, vcloud resources names must not contain any of the following characters: ' : ; * ? , < > / \ | .

Administrator response:

Rename the identified resources to remove unsupported characters from their name. Or, remove them from your backup selection.

GVM1502E **You have selected the vApp from a different organization VDC. For restore tasks, all selected vApps must belong to the same organization VDC. Change your selections and retry the operation.**

GVM1503E **The vApp vApp name exists. Choose a different vApp name to be the target of the restore.**

GVM2001E **Your selection of items to back up has caused the backup task definition to require *count* characters, which exceeds the 512 character limit. Please create multiple backup tasks with less items per task.**

Explanation:

The backup task definition has a 512 character limit, and the total number of characters for the selected items exceeds this limit.

Administrator response:

Create multiple backup tasks with less items per task

GVM2002E **The Organization VDC node can not be included because its Provider VDC node is not included. Please select the include**

checkbox for the Provider VDC node first, and try again.

GVM2004E **The nodename *node name* is already in use. Please uncheck the register node checkbox or choose another nodename.**

Explanation:

The node name chosen already exists on the server. Either choose to not register it or use another name.

Administrator response:

Pick another node name to use. If you want to re-use this existing node, then unselect the 'Register Node' checkbox.

GVM2007E **The Organization VDC node can not be registered because its provider VDC is not valid.**

GVM2008E **The Organization VDC name *OVDC name* is invalid. For information about supported characters, refer to the *IBM Storage Protect Administrator's Reference publication section about naming IBM Storage Protect objects*.**

GVM2009I **This task was skipped because it was not necessary. No further action is required.**

GVM2010W **Internet explorer version *version* is not supported, please use a supported version or another browser. You may see visual and functional issues if you continue to use this unsupported browser.**

Explanation:

Due to differences in Internet Explorer implementation by version number, only specific versions are supported. The use of a standards-compliant browser such as Mozilla Firefox is recommended. However, if you are accessing the GUI as a plug-in from the vSphere Client, you are limited to using the Internet Explorer browser installed on the system where the vSphere client is installed.

Administrator response:

Use a supported version of Internet Explorer or another browser. Supported browser versions are documented in the online help.

GVM2011W **The browser *version* is not supported, please use a supported browser. You may see visual and functional issues if you continue to use this unsupported browser.**

Explanation:

Due to differences in browser implementations, only specific versions are supported.

Administrator response:

Use a supported browser. Supported browser versions are documented in the online help.

GVM2012E **At least one virtual machine that you have selected for restore to alternate location already exists in the Datacenter, so restore is not allowed. To restore to an alternate location when the destination virtual machine already exists, select only one virtual machine for the restore operation and choose a new name for the destination virtual machine. Duplicated VM: *VM name***

Explanation:

When restoring to an alternate location, the destination virtual machine must not already exist.

Administrator response:

Use the single virtual machine restore wizard so that you can rename the destination virtual machine.

GVM2012W **Target datastore not found, select a different destination datastore.**

GVM2013E **The user *User Name* is not authorized to any managed datacenters. Contact your system administrator.**

GVM2014E **You do not have required permissions to view virtual machines for this Event.**

GVM2015E **You do not have required permissions to view restore points for this virtual machine.**

GVM2016E **You do not have required permissions to view some attached points.**

GVM2017E **You do not have required permissions to view restore points for this datastore.**

GVM2018E **You do not have required permissions to detach for the restore point.**

GVM2019E **An error occurred processing user permissions. Contact your system administrator.**

GVM2020I **Some datacenters are not shown due to permissions requirements.**

GVM2021E **You do not have permissions to cancel this task.**

GVM2022I	The task is still in the starting state, please refresh the task and try the cancel again.
-----------------	---

GVM2025E	An error occurred while writing to the flrConfig.props configuration file.
-----------------	---

Explanation:

The flrConfig.props file contains configuration options for file level restore processing. Possible reasons for this error include the following situations: The flrConfig.props file is not in the IBM Tivoli Data Protection for VMware installation directory. The flrConfig.props file is write-protected.

Administrator response:

Verify that the file exists in the IBM Tivoli Data Protection for VMware installation directory and that the file is not write-protected.

GVM2026E	The local mount proxy node pair cannot be removed while the file level restore feature is enabled.
-----------------	---

Explanation:

File level restore processing requires a local mount proxy node.

Administrator response:

Disable the file level restore feature. Then, choose whether you want to remove the mount proxy node pair.

GVM2027E	An error occurred while reading the flrConfig.props configuration file.
-----------------	--

Explanation:

The flrConfig.props file contains configuration options for file level restore processing. The file cannot be read. A common reason for this error is that the file is read-protected.

Administrator response:

Verify that the file is not read-protected.

GVM2028E	The REST API service is unavailable. The web server could not be contacted.
-----------------	--

Explanation:

The web service on the remote machine on which the data mover is installed could not be contacted.

Administrator response

Perform one or more of the following steps to try and determine the problem:

- Verify that the web server service on the remote data mover host is running.
- Verify that this web server is accessible over the network.

GVM2029E	vCenter credentials are incorrect.
-----------------	---

Explanation:

In order to install the data mover services, valid vCenter credentials are needed.

Administrator response:

Verify that you have the correct vCenter credentials. Then try the action again.

GVM2030W	The specified user does not have sufficient permissions to access the following data centers: <i>list of data centers</i>. Click OK to continue or cancel to enter another user name.
-----------------	--

Explanation:

The user credentials that you use to authenticate to the vCenter Server must have the correct privileges to access the VMware datacenters.

Administrator response:

Verify that you have the correct privileges. See the vCenter Server credentials online help page to see the privileges that are required.

GVM2031I	The specified user has sufficient permissions to access the following data centers: <i>list of data centers</i>. Click OK to continue or cancel to enter another user name.
-----------------	--

Explanation:

The user credentials that you use to authenticate to the vCenter Server must have the correct privileges to access the VMware datacenters.

Administrator response:

Verify that you have the correct privileges. See the vCenter Server credentials online help page to see the privileges that are required.

GVM2032W	The specified user does not have sufficient permissions to access any data center. Click OK to continue or cancel to enter another user name.
-----------------	--

Explanation:

The user credentials that you use to authenticate to the vCenter Server must have the correct privileges to access the VMware datacenters.

Administrator response:

Verify that you have the correct privileges. See the vCenter Server credentials online help page to see the privileges that are required.

GVM2033I	Some datacenters are not available because they have the same name for one or more
-----------------	---

datacenters. Datacenters with the same name are not supported.

GVM2132E **An error occurred when connecting to the IBM Tivoli Storage Manager server *server name*. Either your admin ID or password is not valid, or the TCPPORT number was entered in the admin port field instead of the TCPADMINPORT or SSLTCPADMINPORT number.**

Explanation:

See message.

Administrator response:

Launch the Configuration Editor from the Configuration Tab and enter a valid ID or password for your IBM Tivoli Storage Manager Server.

GVM2133E **The password for the administrative user ID *admin id* expired on the IBM Tivoli Storage Manager server *server name*.**

Explanation:

Your IBM Tivoli Storage Manager administrative password has expired.

Administrator response:

Contact your IBM Tivoli Storage Manager Server administrator to reset the password for the administrative user ID.

GVM2134E **The IBM Tivoli Storage Manager server port number *tcp port* is incorrect. The expected value for this port is *tcp port from query*, which is the value of the TCPPORT**

option. Please enter the expected value using the configuration wizard.

Explanation:

The value entered in the IBM Tivoli Storage Manager server port field must match the TCPPORT option on the IBM Tivoli Storage Manager server.

Administrator response:

Use the configuration wizard to change the IBM Tivoli Storage Manager server port field to the correct value.

GVM2135E **This schedule contains an unsupported option so it cannot be edited. This situation can occur when the schedule was created or updated by a tool other than the Data Protection for VMware GUI.**

GVM2136E **An error occurred while processing a VMCLI command, and the GUI session will be closed. Log in and try the operation again. If the problem persists, contact your administrator.**

GVM3000E **Windows domain credentials are incorrect. Open the Configuration Editor, go to File Restore page, and try entering the credentials again.**

Explanation:

The Windows domain credentials that was entered on the File Restore page in the Configuration Editor is incorrect.

Administrator response:

Run the Configuration Editor again and re-enter the correct Windows domain credentials.

Appendix D. IBM Storage Protect recovery agent messages

This information contains explanations and suggested actions for messages issued by the recovery agent.

Beginning with version 8.1, recovery agent messages contain the IBM Storage Protect product name where the product is referenced in the message. This change is not reflected in the documented messages, which still contain the product name Tivoli Storage Manager.

FBP messages

FBP0001E The Recovery Agent is already running.

Explanation:

This issue is encountered when multiple users are logged on to the system and attempting to run the Recovery Agent. Only one active Recovery Agent instance is supported.

System action:

The Recovery Agent was not started.

User response:

To resolve this issue, close the current Recovery Agent instance or start the Recovery Agent on a different system.

FBP0002E The Tivoli Storage Manager server connection cannot be removed.

Explanation:

There are currently active instant restore sessions or mounted volumes that require the connection to the Tivoli Storage Manager server. As a result, the existing connection cannot be removed.

System action:

The remove connection operation is canceled.

User response:

To resolve this issue, wait until the instant restore sessions complete. Or, forcibly end the instant restore sessions or mounted volumes and then disconnect the Tivoli Storage Manager server.

FBP0003E 'Authentication node' and 'Target node' cannot specify the same node.

Explanation:

Three node authentication methods are available to access snapshots on the Tivoli Storage Manager server: 'Asnodename' authenticates with a proxy node, 'Fromnode' authenticates with a node that contains limited access, and 'Direct' authenticates directly. When 'Fromnode' or 'Asnodename' are selected, a target node must be specified. The target node is the Tivoli Storage Manager node where the virtual machine backups are located.

System action:

The system waits for a user response.

User response:

Specify the correct 'Authentication node' and 'Target node'. See information about the node authentication methods in the product documentation.

FBP0004E Recovery Agent failed to mount.

Explanation:

The mount operation on the Windows Recovery Agent proxy host failed.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs for information about why the mount operation failed.

FBP0005E RAID mirror status was not obtained.

Explanation:

During the instant restore session, the Recovery Agent failed to obtain the status of the mdadm mirror device.

System action:

An attempt to recover the instant restore session is being made.

User response:

Check the instant restore status in the Recovery Agent GUI and the Recovery Agent engine logs for solutions to this problem.

FBP0006E Incorrect parameters were specified during the Recovery Agent mount operation.

Explanation:

The mount operation on the Windows Recovery Agent proxy host failed because incorrect parameters were specified.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs for information about why the mount operation failed.

FBP0007E The selected snapshot is already mounted.

Explanation:

The Windows Recovery Agent proxy host identified the selected snapshot as already mounted to the requested target.

System action:

None.

User response:

The instant restore session or mounted snapshot is available for use.

FBP0010E Failed to dismount.

Explanation:

The dismount operation on the Windows Recovery Agent proxy host failed.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs for information about why the dismount operation failed.

FBP0011E The node data was not retrieved.

Explanation:

The Windows Recovery Agent proxy host failed to retrieve the node data when querying the Tivoli Storage Manager server.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs for information about why the querying operation failed.

FBP0012E Tivoli Storage Manager server snapshots were not found.

Explanation:

The Windows Recovery Agent proxy host failed to expose the snapshots on the specified Tivoli Storage Manager server.

System action:

The operation is canceled.

User response:

Verify that the correct Tivoli Storage Manager server and node that own the snapshots are specified.

FBP0013E The snapshot was not found. Click 'Refresh'.

Explanation:

The selected snapshot was not found on the Tivoli Storage Manager server.

System action:

The operation is canceled.

User response:

Click Refresh in the Recovery Agent GUI to load the current snapshots on the Tivoli Storage Manager server.

FBP0016E Invalid parameters were specified.

Explanation:

The mount operation on the Windows Recovery Agent proxy host failed.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs for information about why the mount operation failed.

FBP0017E The mdadm version is not supported.

Explanation:

The installed version of the mdadm utility on the Linux machine is not supported.

System action:

The operation is canceled.

User response:

Upgrade the mdadm utility on your Recovery Agent Linux machine to a supported version. For current software requirements, see [technote 1505139](#).

FBP0018E The mdadm utility was not found.

Explanation:

The mdadm utility is not installed on the Linux machine.

System action:

The operation is canceled.

User response:

Install the mdadm utility on your Recovery Agent Linux machine. For current software requirements, see [technote 1505139](#).

FBP0019E The iscsiadm version is not supported.

Explanation:

The installed version of the iscsiadm utility (for Recovery Agent on RedHat Linux machines) or open-iscsi (for Recovery Agent on SUSE Linux machines) is not supported.

System action:

The operation is canceled.

FBP0020E The iscsiadm utility was not found.

Explanation:

The iscsiadm utility is not installed on the Linux machine.

System action:

The operation is canceled.

User response:

Install the iscsiadm utility on your Recovery Agent Linux machine. For current software requirements, see [technote 1505139](#).

FBP0021E The lsscsi version is not supported.

Explanation:

The installed version of the lsscsi utility is not supported.

System action:

The operation is canceled.

User response:

Upgrade the lsscsi utility on your Recovery Agent Linux machine to a supported version. For current software requirements, see [technote 1505139](#).

FBP0022E The lsscsi utility was not found.

Explanation:

The lsscsi utility is not installed on the Linux machine.

System action:

The operation is canceled.

User response:

Install the lsscsi utility on your Recovery Agent Linux machine. For current software requirements, see [technote 1505139](#).

FBP0023E The Secure Shell (SSH) version is not supported.

Explanation:

The installed version of the SSH client is not supported.

System action:

The operation is canceled.

User response:

Upgrade the SSH client on your Recovery Agent Linux machine to a supported version. For current software requirements, see [technote 1505139](#).

FBP0024E The Secure Shell (SSH) was not found.

Explanation:

The SSH client is not installed on the Linux machine.

System action:

The operation is canceled.

User response:

Install the SSH client on your Recovery Agent Linux machine. For current software requirements, see [technote 1505139](#).

FBP0025E Not all instant restore sessions were stopped.

Explanation:

The Recovery Agent was unable to stop all instant restore sessions.

System action:

Some of the instant restore sessions are still visible in the Recovery Agent 'Instant Restore' panel.

User response:

Try stopping the instant restore sessions one after the other. If the problem persists, check the Linux Recovery Agent engine log file. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf. Also check the Windows Recovery Agent proxy host logs for any errors related to this issue.

FBP0026E Failed to read instant restore session.

Explanation:

The Recovery Agent was unable to read the instant restore status file.

System action:

Information about the instant restore is not available.

User response:

Try restarting the instant restore session. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0027E Mount session already exists.

Explanation:

The selected snapshot is already mounted to the requested target.

System action:

None.

User response:

The instant restore session or mounted snapshot is available for use.

FBP0028E Failed to create the mount sessions directory.

Explanation:

The Recovery Agent was unable to create the directory for the mount operation.

System action:

The operation is canceled.

User response:

Try the mount operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0029E Failed to encrypt node credentials.

Explanation:

The Recovery Agent was unable to encrypt the node credentials.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0030E Failed to decrypt node credentials.

Explanation:

The Recovery Agent was unable to decrypt the node credentials.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0031E Failed to remove mount session.

Explanation:

The Recovery agent was unable to delete the mount status file.

System action:

The mounted snapshot remain in the Recovery Agent 'Mounted Volumes' panel.

User response:

Try unmounting the mounted snapshot again. If the problem persists, check the Linux Recovery Agent engine log file. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf for more information. Also check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0032E Instant restore session already exists.

Explanation:

A similar instant restore session already exists or a similar instant restore status file exists.

System action:

The new instant restore session is canceled.

User response:

Check the Linux Recovery Agent engine log file for more information. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0033E Failed to create the instant restore sessions directory.

Explanation:

The Recovery Agent was unable to create the directory for the instant restore operation.

System action:

The instant restore operation is canceled.

User response:

Try the instant restore operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0034E Failed to remove the instant restore session.

Explanation:

The Recovery Agent was unable to delete the instant restore status file.

System action:

The instant restore session remains in the Recovery Agent 'Instant Restore' panel.

User response:

Try stopping the instant restore session again. If the problem persists, check the Linux Recovery Agent engine log file. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf for more information. Also check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0035E Failed to read from the configuration file that is used for mount and instant restore operations.

Explanation:

The Recovery Agent was unable to read the configuration file.

System action:

Information about the mount or instant restore is not available.

User response:

Try the operation again. If the problem persists check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0036E Failed to write to the configuration file that is used for mount and instant restore operations.

Explanation:

The Recovery Agent was unable to write to the configuration file.

System action:

Information about the mount or instant restore is not available.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0037E **Failed to read from the configuration file section that is used for mount and instant restore operations.**

Explanation:

The Recovery Agent was unable to read the configuration file.

System action:

Information about the mount or instant restore is not available.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0038E **Failed to write to the configuration file section that is used for mount and instant restore operations.**

Explanation:

The Recovery Agent was unable to write to the configuration file.

System action:

Information about the mount or instant restore is not available.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0039E **Failed to unmount. Device is busy.**

Explanation:

The Recovery Agent was not able to unmount the file system of the selected mounted snapshot because the file system is in use.

System action:

The unmount operation is canceled.

User response:

Close any application that might be accessing this volume. Then, try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0040E **Not all mount sessions were unmounted.**

Explanation:

The Recovery Agent did not unmount all mounted snapshots.

System action:

Some of the mounted snapshots sessions are still visible in the Recovery Agent 'Mounted Volumes' panel.

User response:

Try to unmount the mounted snapshots one after the other. If the problem persists check the Linux Recovery Agent engine log file. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0041E **Failed to retrieve data from the Recovery Agent CLI.**

Explanation:

The Recovery Agent was unable to read the Recovery Agent CLI (TDPVMWareShell) output file.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0042E **Failed to parse data from the Recovery Agent CLI.**

Explanation:

The Recovery Agent was unable to parse the data from the Recovery Agent CLI (TDPVMWareShell) output file.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux Recovery Agent engine log file. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0043E **Failed to create query for the Recovery Agent CLI.**

Explanation:

The Recovery Agent was unable to create the Recovery Agent CLI (TDPVMWareShell) output file.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0044E **Failed to retrieve mount data from the Recovery Agent CLI.**

Explanation:

The Recovery Agent was unable to create the Recovery Agent CLI (TDPVMWareShell) output file.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0045E Failed to create mount query for the Recovery Agent CLI.

Explanation:

None.

FBP0046E Failed to connect to the Recovery Agent CLI.

Explanation:

The Linux Recovery Agent was unable to retrieve the Recovery Agent CLI (TDPVMWareShell) installation path from the registry on the Windows machine.

System action:

The operation is canceled.

User response:

Check the connectivity to the Windows machine, verify that the SSH is configured correctly, and that the user defined on Cygwin has administrative privileges. For more information, see the product documentation.

FBP0047E Failed to create the mount directory.

Explanation:

The Recovery Agent was unable to locate or create the directory for the mount operation.

System action:

The mount operation is canceled.

User response:

Try the mount operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0048E Failed to mount the file system of the snapshot.

Explanation:

The Recovery Agent succeeded connecting to the mounted snapshot iSCSI device. However, the Recovery Agent was unable to mount the file system of the snapshot.

System action:

The mount operation failed. The Recovery Agent automatically attempts to mount the file system every 5 minutes.

User response:

Since the mounted snapshot is available as an iSCSI device, attempt to mount the file system of the device. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0049E Failed to set the SSH server address.

Explanation:

The Recovery Agent was unable to set the SSH server address. The specified server address might be incorrect.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux Recovery Agent engine log file for more information. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0050E Failed to set SSH user name.

Explanation:

The Recovery Agent cannot access the Windows Recovery Agent CLI (TDPVMWareShell) using the defined 'SSH login' user.

System action:

The operation is canceled.

User response:

Using Secure Shell verify that you can connect to the Windows Recovery Agent CLI machine using the user defined under 'Settings'>'SSH login'.

FBP0051E Failed to run SSH command.

Explanation:

None.

FBP0052E Failed to send query to the TDPVMware Shell.

Explanation:

None.

FBP0053E Failed to send query to the Recovery Agent CLI.

Explanation:

While using SSH, the Recovery Agent failed to send an input command file to the Windows Recovery Agent CLI (TDPVMWare Shell). This issue might be caused by an SSH user without read and write privileges on the Windows Recovery Agent CLI machine.

System action:

The operation is canceled.

User response:

Verify that the SSH user defined under 'Settings'>'SSH login' has read and write privileges on the Windows Recovery Agent CLI machine. Also check the Linux Recovery Agent engine log file for more information. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0056E Failed to locate the block device specified for the mount point.

Explanation:

The Recovery Agent failed to locate the block device for the mount point that was specified for the instant restore operation.

System action:

The instant restore operation is canceled.

User response:

Verify that the mount point specified for the instant restore operation is correct.

FBP0057E Failed to locate the mount point for the specified block device.

Explanation:

The Recovery Agent failed to locate the mount point for the specified instant restore block device.

System action:

The instant restore operation is canceled.

User response:

Verify that the specified block device for the instant restore operation is correct and mounted.

FBP0058E The specified mount point is not mounted on the block device.

Explanation:

The specified mount point is not mounted on the specified block device.

System action:

The instant restore operation is canceled.

User response:

Verify that the specified mount point and block device are correct, and that the mount point is mounted on that block device.

FBP0059E Restore is not allowed to '/' or '/boot'.

Explanation:

The Recovery Agent does not support instant restore operations to target devices that are '/' or '/boot' volumes.

System action:

The instant restore operation is canceled.

User response:

Specify a different target device for the instant restore operation.

FBP0060E Restore is not allowed to RAID devices.

Explanation:

The Recovery Agent does not support instant restore operations to RAID devices.

System action:

The instant restore operation is canceled.

User response:

Specify a different target device for the instant restore operation.

FBP0061E The restore operation failed to start.

Explanation:

The Recovery Agent failed to start the instant restore operation.

System action:

The instant restore operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux Recovery Agent engine log file for more information. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0062E Failed to get the iSCSI initiator name.

Explanation:

No iSCSI initiator was specified.

System action:

The operation is canceled.

User response:

Verify that the iSCSI initiator name is specified correctly.

FBP0063E iSCSI target is not logged in.

Explanation:

None.

FBP0064E Failed to resolve the IP address to a hostname.

Explanation:

The Recovery Agent failed to associate a hostname with the specified IP address.

System action:

The operation is canceled.

User response:

Verify that the IP or hostname are correct. Then, try the operation again. If the problem persists, check the Linux Recovery Agent engine log file for more information. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0065E The iSCSI service was not found.

Explanation:

The iSCSI daemon is not installed on the Recovery Agent Linux machine.

System action:

The operation is canceled.

User response:

Install the required iSCSI daemon on the Recovery Agent Linux machine. See information about related iSCSI tasks in the product documentation.

FBP0066E Failed to start the iSCSI daemon.

Explanation:

The Recovery Agent was unable to start the iSCSI daemon.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0067E Failed to discover iSCSI targets.

Explanation:

The Recovery Agent was unable to discover the iSCSI targets.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0068E Failed to log in to the iSCSI target.

Explanation:

The Recovery Agent was unable to log in to the iSCSI target.

System action:

The operation is canceled.

User response:

Check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0069E Failed to log out of the iSCSI target.

Explanation:

The Recovery Agent was unable to log out of the iSCSI target.

System action:

The operation is canceled.

User response:

Check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0070E Failed to delete the iSCSI target.

Explanation:

The Recovery Agent was unable to delete the iSCSI target.

System action:

The operation is canceled.

User response:

Check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0071E Failed to identify the block device for the iSCSI target.

Explanation:

The Recovery Agent successfully logged in to the iSCSI target; however, the iSCSI block device was not found.

System action:

The operation is canceled.

User response:

Check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0072E Failed to create the RAID mirror device.

Explanation:

The Recovery Agent was unable to create the RAID mirror device using the mdadm utility.

System action:

The instant restore operation is canceled.

User response:

Check the Linux system log (usually located at /var/log/messages) for any errors related to this issue.

FBP0073E Failed to stop the RAID mirror device.

Explanation:

The Recovery Agent uses the mdadm utility to stop the RAID mirror device. However, the Recovery Agent was unable to stop the RAID mirror device.

System action:

The instant restore operation is canceled.

User response:

Check the Recovery Agent engine log file on the Linux machine for more details about why the device did not stop. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf. Also check the Linux system log (usually located at /var/log/messages).

FBP0074E Failed to add the target block device to the RAID mirror.

Explanation:

The Recovery Agent uses the mdadm utility to add the target block device to the RAID mirror. However, the Recovery Agent was unable to add the target block device to the RAID mirror.

System action:

The instant restore operation is canceled.

User response:

Check the Recovery Agent engine log file on the Linux machine for more details about why the device was unable to add the target block device to the RAID mirror. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf. Also check the Linux system log (usually located at /var/log/messages).

FBP0075E Failed to mark the target block device in RAID mirror as faulty.

Explanation:

The Recovery Agent uses the mdadm utility to mark the target block device in the RAID mirror. However, the Recovery Agent was unable to mark the target block device as faulty.

System action:

The instant restore session is paused.

User response:

Check the Recovery Agent engine log file on the Linux machine for more details about why the device was not marked. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf. Also check the Linux system log (usually located at /var/log/messages).

FBP0076E Failed to remove the target block device from the RAID mirror.

Explanation:

The Recovery Agent uses the mdadm utility to remove the target block device from the RAID mirror. However, the Recovery Agent was unable to remove the target block device.

System action:

The instant restore operation is canceled.

User response:

Check the Recovery Agent engine log file on the Linux machine for more details about why the device was not removed. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf. Also check the Linux system log (usually located at /var/log/messages).

FBP0079E Unknown key.

Explanation:

None.

FBP0080E The operation timed-out.

Explanation:

The Recovery Agent engine scripts did not reply to the Recovery Agent GUI in a timely manner.

System action:

The operation is canceled.

User response:

Try the operation again. If the problem persists, check the Recovery Agent engine log file on the Linux machine. This file is usually located at /opt/tivoli/tsm/TDPVMWare/mount/engine/var/TSM4VE_IR_LOG_0040.sf.

FBP0081E Internal error.

Explanation:

None.

FBP0083E The snapshot does not contain a partition with a supported file system.

Explanation:

The Recovery Agent successfully parsed the partition structure of the disk. However, the partitions do not use any of these supported file systems: FAT, NTFS, EXT2, EXT3, EXT4, or ReiserFS. As a result, volume-level operations, such as 'Instant Restore' and 'Mount as Virtual Volume', are not supported for this snapshot.

System action:

The operation is canceled.

User response:

To restore data from the selected snapshot, use the Windows Recovery Agent proxy host to mount and expose the snapshot. See information about mounting as iSCSI targets in the product documentation.

FBP0084E Failed to retrieve partitions.

Explanation:

The Windows Recovery Agent proxy host was unable to retrieve the partition list from the disk snapshot.

System action:

The operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs about why the partition list was not retrieved.

FBP0085E Recovery Agent can connect only to a Tivoli Storage Manager server node.

Explanation:

None.

FBP0086E Failed to remove the Tivoli Storage Manager server connection.

Explanation:

The Windows Recovery Agent proxy host reported that there are active instant restore sessions or mounted volumes that require the connection to the Tivoli Storage Manager server. As a result, the existing connection cannot be removed.

System action:

The remove connection operation is canceled.

User response:

Check the Windows Recovery Agent proxy host logs about the active instant restore sessions or mounted volumes.

FBP0088E **Mount operation failed because the Write cache is either full or configured incorrectly.**

Explanation:

Mount and instant restore sessions (that run on the Linux machine) access the Virtual Volume write cache on the Windows Recovery Agent proxy host. This proxy host reported that the write cache is unavailable.

System action:

The operation is canceled.

User response:

Check the log files on the Windows Recovery Agent proxy host for information about why the write cache is unavailable. Verify that the Virtual Volume write cache is configured correctly in the Windows Recovery Agent GUI. See information about setting the Virtual Volume write cache options in the product documentation.

FBP0089E **The Recovery Agent GUI storage type option is 'Tape', and the requested media is busy.**

Explanation:

When the Recovery Agent GUI storage type option specifies 'Tape', only a single snapshot can be mounted.

System action:

The operation is canceled.

User response:

Dismount the currently mounted snapshot before you attempt to mount another snapshot.

FBP0090E **Operation timed-out. Manual intervention might be required.**

Explanation:

The Linux Recovery Agent GUI operation timed-out during a mount or instant restore operation.

System action:

The operation is canceled.

User response:

Follow the "Responding to a timeout during a file restore or an instant restore (Linux)" procedure

documented in the Data Protection for VMware Installation and User's Guide.

FBP0091E **The selected disk is not an MBR disk.**

Explanation:

The Recovery Agent cannot parse the partition structure of the disk, because the disk is not a Basic, MBR-based disk. Volume-level operations, such as 'Instant Restore' and 'Mount as Virtual Volume', are not supported for this snapshot.

System action:

Mount and instant restore operations are disabled.

User response:

To restore data from the selected snapshot, use the Windows Recovery Agent proxy host to mount and expose the snapshot using 'Mount as iSCSI target' method.

FBP0092E **Operation failed. Recovery Agent is initializing.**

Explanation:

Mount and instant restore operations cannot proceed when the Windows Recovery Agent proxy host is initializing.

System action:

The mount or instant restore operation is canceled.

User response:

Wait until the Windows Recovery Agent proxy host completes initializing. Then, try the operation again.

FBP0093E **The mount point is already in use by another mount session.**

Explanation:

The mount operation failed because the target mount point is currently in use.

System action:

The mount operation is canceled.

User response:

Specify a target mount point that is not in use.

FBP0094E **The mount point is already in use by another instant restore session.**

Explanation:

The instant restore operation failed because the target mount point is currently in use.

System action:

The instant restore operation is canceled.

User response:

Specify a target mount point that is not in use.

FBP1001I **The folder: *folder path* does not exist. Do you want to create the folder?**

Explanation

The selected folder does not exist. You can create the folder by proceeding with the instructions displayed on the screen.

System action

The system waits for a user response.

User response

Specify Yes to create the folder and continue with the operation. Specify No to not create the folder and end the current operation.

FBP1003I Specify a folder for the cache files

Explanation

A folder that is used to store the virtual disk and virtual volume cache files was not specified. The cache is used to store write operations to the virtual disk and virtual volume during the mount operation.

System action

The system waits for a user response.

User response

Specify a folder to store the virtual disk and virtual volume cache files.

FBP1005I *mount path* already has a mounted virtual volume. Do you want to dismount the volume before you continue?

Explanation

The selected path already has a mounted virtual volume. This existing mounted virtual volume must be dismounted to continue with the current mount operation.

System action

The system waits for a user response.

User response

Specify Yes to dismount the existing mounted virtual volume and continue with the current mount operation. Specify No to end the current mount operation.

FBP1009I Recovery Agent cannot be closed while mounted volumes exist. Do you want to dismount all volumes?

Explanation

Existing mounted virtual volumes and virtual disks must be dismounted before closing the Recovery Agent application.

System action

The system waits for a user response.

User response

Specify Yes to dismount all the existing mounted virtual volumes and virtual disks and close the Recovery Agent application. Specify No to not dismount all the existing mounted virtual volumes and virtual disks and return to the Recovery Agent application.

FBP1008I Recovery Agent cannot be started because it is being used by another user

Explanation

The Recovery Agent application is already running on this machine. For example, this situation occurs when the application is started by another user on another session.

System action

The operation is canceled.

User response

Log in to the machine as the user that started the Recovery Agent application.

FBP1011I The Virtual Volume Driver was registered successfully

Explanation

Successful Recovery Agent Virtual Volume Driver registration is required to use the mount function.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP1013I **Connection to Tivoli Storage Manager server resumed.**

Explanation

Communication between the Recovery Agent application and the Tivoli Storage Manager server is established.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP1014I **You must restart the Recovery Agent application for changes to the Data Access options to take effect.**

Explanation

Recent changes to the Data Access options cannot be implemented until the Recovery Agent application is restarted.

System action

The Recovery Agent application operates with the existing Data Access options setting until it is restarted.

User response

Restart the Recovery Agent application so that changes to the Data Access options are implemented.

FBP1100I **Received CONTINUE from service manager**

Explanation

The Recovery Agent service received a SERVICE_CONTROL_CONTINUE command.

System action

The Recovery Agent service resumes activities.&msgnl;The latest Recovery Agent driver events are written to the Recovery Agent log file.

User response

The Recovery Agent application is ready for operations.

FBP1101I **Received PAUSE from service manager**

Explanation

The Recovery Agent service received a SERVICE_CONTROL_PAUSE command.

System action

This command has no effect on the Recovery Agent service.

User response

The Recovery Agent application is ready for operations.

FBP1102I **Received a STOP command from the service manager**

Explanation

The Recovery Agent service received a SERVICE_CONTROL_STOP command. This message is issued when the service is stopped from the service manager. Opening the Recovery Agent GUI from the Start menu sends a stop command to the service.

System action

The Recovery Agent service stops.

User response

The Recovery Agent application is ready for operations.

FBP1103I **Service stopped, reporting to service manager**

Explanation

The Recovery Agent service stops. The service can be restarted from the service manager.

System action

The Recovery Agent application stops.

User response

No user action required.

FBP1104I **--- Recovery Agent version string started ---**

Explanation

The Recovery Agent application started. The version string is printed in the message.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP1300I **License is OK.**

Explanation

The license file is valid.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP1301W **Not for resale license.**

Explanation

The license file is valid.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP1302E **Error accessing license file.**

Explanation

The license file could not be accessed. This situation occurs when the license file cannot be located, it cannot be opened because of permission restrictions, or the file is corrupted.

System action

The Recovery Agent application stops.

User response

Obtain a new license for the Recovery Agent application.

FBP1303E **Corrupted license file.**

Explanation

The license registration string is not valid.

System action

The Recovery Agent application stops.

User response

Obtain a new license for the Recovery Agent application.

FBP1304E **Trial period has expired.**

Explanation

The license registration string is not valid.

System action

The Recovery Agent application stops.

User response

Obtain a new license for the Recovery Agent application.

FBP1305I **Try and buy license. days left.**

Explanation

The license file is valid.

System action

Operation completed successfully.

User response

The Recovery Agent application is ready for operations.

FBP5003W ***target volume : The repository is not loaded.Load the repository and resume the session.***

Explanation

The instant restore session is unable to resume since the Recovery Agent failed to reestablish connection to the Tivoli Storage Manager server where the

backup snapshots are located. The restored volume is inaccessible while the session is paused.

System action

The instant restore session is paused.

User response

Click "Select IBM Storage Protect server" in the Recovery Agent GUI to connect to the Tivoli Storage Manager server and resume the instant restore session.

FBP5005W	Windows indicates the destination volume <i>target volume</i> might be a network-mapped drive. If volume <i>target volume</i> again is confirmed as a network-mapped drive, the operation fails. Continue anyway?
-----------------	--

Explanation

The Recovery Agent does not support instant restore sessions to a network-mapped drive.

System action

The instant restore session is canceled.

User response

Specify a destination volume that is not on a network-mapped drive.

FBP5007W	Read block from the source has failedThe problem might have been caused by a network failure. See log file for more details. If the problem was caused by a network failure, correct the problem and resume the session.
-----------------	---

Explanation

The instant restore session is unable to retrieve data from the Tivoli Storage Manager server. As a result, the instant restore session is paused. The problem might be caused by a network failure. The restored volume is inaccessible while the session is paused.

System action

The instant restore session is paused.

User response

Check the Recovery Agent logs for information regarding the cause of the problem. After resolving the issue, resume the session.

FBP5008W	There are open handles to the volume being restored (<i>volume name</i>). Close any application (such as Windows Explorer or a command prompt) that might be accessing this volume and try again. If you select Ignore, applications using these handles might become unstable once the restore process begins.
-----------------	--

Explanation

The target volume for the restore is in use. Restoring a volume to a viewable storage volume involves overwriting data on that existing storage volume. After the restore begins, the current volume contents are permanently erased.

System action

The system waits for a user response.

User response

Close any application (such as Windows Explorer or a command prompt) that might be accessing this volume and try the operation again. If you select Ignore, applications that currently use these open handles might become unstable when the restore operation begins.

FBP5010W	System is low on memory. Write operations to virtual volumes might be lost.
-----------------	--

Explanation

Changes that are done on mounted volumes are written to memory. As a result, the Recovery Agent can use a large amount of RAM when it operates in read/write mode.

System action

Write operations to virtual volumes might be lost.

User response

Dismount some of the mounted volumes (when possible) or mount the volumes as read-only. The value of the 'Read Ahead cache size' option affects the

memory usage. See information about setting this option in the product documentation.

FBP5011W	Recovery Agent still has <i>number of mounted volumes</i> volumes mounted. Stopping the Recovery Agent might cause the system to become unstable. Are you sure you want to stop the Recovery Agent?
-----------------	--

Explanation

Stopping the Recovery Agent without first dismounting the virtual volumes might cause both the system and active applications to become unstable.

System action

The system waits for a user response.

User response

Specify No to prevent the Recovery Agent from stopping, then dismount any mounted volumes. Specify Yes to stop the Recovery Agent, even though both the system and active applications might become unstable.

FBP5012W	The Recovery Agent still has <i>number of active sessions</i> active instant restore session. These sessions will be paused and the restored volumes will appear unformatted, until the Recovery Agent service restarts. Are you sure you want to stop the Recovery Agent?
-----------------	---

Explanation

Stopping the Recovery Agent without waiting for the restore to complete makes the restored volumes display as if they are unformatted. The restore process resumes when the Recovery Agent service restarts.

System action

The system waits for a user response.

User response

Specify No to prevent the Recovery Agent from stopping and to continue the instant restore session. Specify Yes to stop the Recovery Agent, even though the volumes that are still being processed display as if they are unformatted. Do not attempt to format these volumes as such an attempt causes data loss.

FBP5013W	Abort selected sessions? All data is lost and volumes require reformatting.
-----------------	--

Explanation

Aborting the instant restore sessions causes the loss of all data that was written to the restored volumes. The restored volumes display as unformatted and require reformatting.

System action

The system waits for a user response.

User response

Specify Yes to abort the instant restore sessions. All data that was written to the restored volumes during these sessions is lost. Specify No to continue the instant restore sessions.

FBP5015W	All data on target drive <i>volume name</i> will be lost. Note 1: Successful instant restore processing requires sufficient network connectivity and bandwidth to the repository. Note 2: Use of instant restore is recommended only for applications that primarily issue READ I/O's. Do you want to continue?
-----------------	--

Explanation

Instant restore processing overwrites data on the target storage volume. A sufficient data transfer rate from the Tivoli Storage Manager server is required for a successful instant restore operation.

System action

The system waits for a user response.

User response

Click YES to confirm that you understand the effects and to start the instant restore operation.

FBP5017W	There are open files on the volume mounted on: <i>mounted volume name</i>. A forced dismount invalidates all of the open files. Are you sure you want to continue?
-----------------	---

Explanation

The virtual volume that is being dismounted is in use by another application. For example, the volume might be open in Windows Explorer. For virtual iSCSI devices, the iSCSI initiator is still logged on to the device. Forcing this volume to dismount might cause the files or applications that are accessing the volume to become unstable.

System action

The system waits for a user response.

User response

Identify and close any files or applications that are accessing the volume. Or, specify 'Continue' to ignore the warning message and continue dismounting the volume. For iSCSI devices, make sure that the iSCSI initiator is logged off the device.

FBP5018W	The volume selected for restore is located on a clustered disk. See the product documentation for guidelines when restoring clustered volumes. Failure to follow these guidelines might result in data loss, if a hardware or Windows error occurs. Do you want to continue?
-----------------	---

Explanation

Instant restore of a volume in a clustered environment is supported. Other volumes in the cluster are not affected. You can work with the cluster and with the restored volume in parallel. During the instant restore operation, the disk that is being restored cannot fail over if the node fails.

System action

The system waits for a user response.

User response

Specify YES to confirm that you understand the guidelines and to start the instant restore operation.

FBP5020W	The Virtual Volume Driver is not yet registered. Recovery Agent can register the driver now. During registration, a Microsoft Windows Logo warning may be displayed. Accept this warning to allow the registration to complete. Do you want to register the Virtual Volume Driver now?
-----------------	---

Explanation

User should register the Virtual Volume Driver in order to work with the Recovery Agent. This message is displayed following the first attempt to mount after a silent install, since the driver registration is not performed in silent install.

System action

The system waits for a user response.

User response

Specifying Yes will start the registration process. After registration the Recovery Agent application is ready for operations.

FBP5021W	The mounted volume '<i>mounted volume name</i>' is in use. Dismounting the volume might cause the application that is using it to become unstable. Do you want to continue?
-----------------	--

Explanation

The virtual volume that is being dismounted is in use by another application. For example, the volume might be open in Windows Explorer. For virtual iSCSI devices, the iSCSI initiator is still logged on to the device.

System action

The system waits for a user response.

User response

Identify and close any applications that are accessing the volume. Or, specify 'Continue' to ignore the warning message and continue dismounting the volume. For iSCSI devices, make sure that the iSCSI initiator is logged off the device.

FBP5023W	There are active instant restore sessions. These sessions will be paused and the restored volumes will appear unformatted, until the application finishes loading and resumes the sessions. Do you want to continue?
-----------------	---

Explanation

Opening the Recovery Agent UI from the 'Start>All Programs' menu stops the service. The active instant restore sessions are paused until the application finishes loading and resumes the sessions.

System action

The system waits for a user response.

User response

Specify Yes to stop the service and load the Recovery Agent UI. This action pauses and then resumes the instant restore sessions. Specify No to not load the Recovery Agent UI. This action leaves the instant restore sessions to run in the background in service mode.

FBP5025W	There are '<i>num active sessions</i>' sessions using the write cache. Updates to the cache settings are processed when there are no active sessions that use the cache. Do you want to continue?
-----------------	--

Explanation

Updates to the write cache settings were detected. These updates are applied when there are no active sessions that use the cache.

System action

The Recovery Agent application operates with the existing write cache settings.

User response

Update the write cache settings when there are no active sessions that use the cache.

FBP5026W	The size of the write cache is <i>num write cache percentage</i>% full.
-----------------	--

Explanation

Mount and instant restore sessions that run on a Linux machine use the Recovery Agent Virtual Volume write cache for write operations. The Cache size is approaching its maximum limit. Linux mount and instant restore sessions might fail when the cache size reaches its limit.

System action

None.

User response

Do not start a new Recovery Agent mount or instant restore session on the Linux machine until the Virtual Volume write 'Cache size' value decreases. See information about setting this option in the product documentation.

FBP5028W	The file system of the selected partition (<i>partition format</i>) might not be supported by the current operating system. The appropriate File System driver must be installed for Windows to read the volume. Do you want to continue?
-----------------	--

Explanation

The partition file system must be supported and recognized by the Windows operating system where the volume is mounted. This condition is required to view the file structure of the mounted partition. It is recommended to mount volumes with native Linux file systems on a Linux machine.

System action

The system waits for a user response.

User response

Make sure the appropriate file system driver is installed on the Windows operating system where the volume is mounted.

FBP5029W	The connection to Tivoli Storage Manager server was lost.
-----------------	--

Explanation

The connection between the Recovery Agent application and the Tivoli Storage Manager server was lost.

System action

Mounted volumes might become inaccessible. The instant restore sessions are paused.

User response

Check the Recovery Agent logs for information regarding the connection failure. After resolving the issue, resume any paused instant restore sessions.

FBP5030W	No snapshots are available for the selected virtual machine. Or, the Authentication node is not authorized to restore this virtual machine.
-----------------	--

Explanation

No snapshots were located for the selected virtual machine. Either no snapshot completed successfully, or if the 'From node' access method was used, the

Authenticate node does not have permission to restore the selected virtual machine.

System action

The operation is canceled.

User response

If the 'From node' access method is used, make sure sufficient permissions are set for the Authentication node. See the product documentation for details and an example of how to set permissions by using the IBM Storage Protect Data Mover 'set access' command.

FBP5031W	Some snapshots are currently mounted. If you continue, these snapshots will be dismounted. If a mounted volume is currently being used by an application, the application might become unstable. Do you want to continue?
-----------------	--

Explanation

Opening the Recovery Agent UI from the 'Start>All Programs' menu stops the service. The active mounted volumes are dismounted.

System action

The system waits for a user response.

User response

Close any application (such as Windows Explorer or a command prompt) that might be accessing the mounted volumes. Then, open the Recovery Agent UI. If you continue without closing the applications that are accessing the mounted volumes, these applications might become unstable.

FBP5032W	The selected snapshot will not be protected from expiration during this operation. See the product documentation for information about expiration.
-----------------	---

Explanation

When the 'From node' authentication method is used, the Authentication node is granted Read-only access to the target node with the 'set access' command. The target node owns the snapshot. As a result, the snapshot cannot be marked as being in use on the server. Therefore, the snapshot might expire while the restore operation is in progress.

System action

The system waits for a user response.

User response

If you proceed, disable the expiration process for the snapshot during the operation. Or, make sure that no snapshots are created for the restored machine during the restore operation. Otherwise, cancel the restore operation.

FBP5033W	No snapshots exist in the selected node.
-----------------	---

Explanation

Either no snapshot was completed, or the selected Tivoli Storage Manager node is not the node that owns the snapshots.

System action

No snapshot is shown.

User response

If Tivoli Storage Manager for Virtual Environments snapshots were completed, select the Tivoli Storage Manager node that owns the snapshots.

FBP5034W	The snapshots are not protected from expiration during the mount operation. An expiration can produce unexpected results and negatively impact the mount point.
-----------------	--

Explanation

The mounted snapshots will not be marked as being in use on the server. Therefore, the snapshots might expire while the restore operation is in progress.

System action

The system waits for a user response.

User response

If you proceed, make sure that no snapshots are created for the restored machines during the restore operations. Otherwise, enable the expiration protection.

FBP5035W	The selected snapshot is not protected from expiration during this mount operation. An expiration can produce
-----------------	--

unexpected results and negatively impact the mount point.

Explanation

The mounted snapshot will not be marked as being in use on the server. Therefore, the snapshot might expire while the restore operation is in progress.

System action

The system waits for a user response.

User response

If you proceed, make sure that no snapshots are created for the restored machine during the restore operation. Otherwise, enable the expiration protection on the 'settings' panel and perform the operation again.

FBP7003E **The folder:*folder name* is invalid.**

Explanation

The specified path is not a valid folder path.

System action

The system waits for a user response.

User response

Specify a valid folder path.

FBP7004E **The folder:*folder name* could not be created.**

Explanation

The system failed to create the requested folder.

System action

The operation is canceled.

User response

Check the Recovery Agent logs for information about why the folder was not created.

FBP7005E ***folder name* is not empty. Only empty folders can be used as a mount point.**

Explanation

The system can mount a volume snapshot only to an empty folder.

System action

The operation is canceled.

User response

Specify an empty folder path.

FBP7006E **snapshot size (*snapshot size*) is larger than target size (*partition name*)**

Explanation

The volume size of the destination location must be equal to, or greater than, the size of the original volume on the snapshot to be restored.

System action

The instant restore operation is canceled.

User response

Specify a target volume with a size equal to, or greater than, the source volume on the snapshot to be restored.

FBP7007E **A valid iSCSI target name must be specified. Valid iSCSI names consist of the following items: lower-case English characters, digits, '.', ':', and '-'.**

Explanation

When a snapshot is exposed as an iSCSI target, a valid iSCSI target name must be entered.

System action

The system waits for a user response.

User response

Specify a valid iSCSI target name.

FBP7008E **A valid iSCSI initiator name must be specified. Valid iSCSI names consist of the following items: lower-case English characters, digits, '.', ':', and '-'.**

Explanation

The specified iSCSI initiator is not a valid initiator name. When a snapshot is exposed as an iSCSI target, a valid initiator name must be specified by the user.

System action

The system waits for a user response.

User response

Specify a valid iSCSI initiator name.

FBP7009E	A valid folder name must be specified.
-----------------	---

Explanation

The path to the folder is invalid or was not specified.
The path to the mount point for the volume snapshot must include an empty folder.

System action

The system waits for a user response.

User response

Specify a valid path to an empty folder.

FBP7012E	Already connected to a Tivoli Storage Manager server. To connect to a different server, or to a different node within the server, select the current server from the list and click 'Remove'.
-----------------	--

Explanation

Recovery Agent can connect only to a single Tivoli Storage Manager server and node.

System action

The operation is canceled.

User response

Remove the existing connection by selecting the server from the list and click 'Remove'. You cannot remove a connection to a server that has active mounted volumes or instant restore sessions.

FBP7013E	Instant restore is not supported in Tape Mode.
-----------------	---

Explanation

Instant restore of snapshot data that is stored on tape is not supported.

System action

The operation is canceled.

User response

Either migrate the snapshot data to a disk storage pool, or use a different restore method.

FBP7014E	Cannot mount more than one snapshot in Tape Mode.
-----------------	--

Explanation

Only a single snapshot can be mounted when the snapshot data is stored on a tape.

System action

The operation is canceled.

User response

Dismount the currently mounted snapshot before you attempt to mount another snapshot.

FBP7015E	No local volume is available as a destination for instant restore.
-----------------	---

Explanation

Instant restore is done to a local volume that has a volume letter and is not the system volume.

System action

The operation is canceled.

User response

Make sure you have a local volume that has a volume letter and that is not the system volume as a destination for the instant restore.

FBP7016E	The Recovery Agent 'Read Ahead size' option requires a value of 0 through 8192.
-----------------	--

Explanation

The Recovery Agent 'Read Ahead size' value specifies the number of extra data blocks retrieved from the storage device after a read request is sent to a single block.

System action

The system waits for a user response.

User response

Specify a valid 'Read Ahead size' value of 0 through 8192. See information about setting this option in the product documentation.

FBP7017E **The Recovery Agent 'Read Ahead cache size' option requires a value of 1000 through 75000. The value must also be at least 1 block larger than the value of the 'Read Ahead size' option.**

Explanation

Recovery Agent 'Read Ahead cache size' value specifies the size of the cache where the 'Read Ahead' extra data blocks are stored.

System action

The system waits for a user response.

User response

Specify a valid 'Read Ahead cache size' value of 1000 through 75000 and at least 1 block larger than the value of the 'Read Ahead size' option. See information about setting this option in the product documentation.

FBP7018E **The Recovery Agent cannot read the snapshot data from the Tivoli Storage Manager server. Make sure that the storage type setting matches the actual storage device. For example, if the data resides on tape, make sure that the storage type is set to 'Tape'. If this does not solve the problem, check the Tivoli Storage Manager server activity log for additional errors.**

Explanation

The required data does not exist on the server or the data is inaccessible. Inaccessible data is typically caused by a 'Storage type' configuration that does not match the actual storage where the data is stored. Recovery Agent cannot read data stored on Tape or VTL if storage type is set to 'Disk'.

System action

The operation is canceled.

User response

Click 'Settings' in the Recovery Agent GUI and select the correct storage device from which to mount the snapshot. You can select 'Disk/File', 'Tape', or 'VTL'. When the storage type is changed, you must restart the Recovery Agent for the changes to take effect. Also, check the Recovery Agent logs and the Tivoli

Storage Manager server activity log for any additional errors.

FBP7019E **The selected disk is not a basic disk with an MBR partition style.**

Explanation

Recovery Agent could not parse the partition structure of the disk, because the disk is not a Basic, MBR-based disk. Volume-level operations, such as 'Instant Restore' and 'Mount as Virtual Volume', are not supported for this snapshot.

System action

Volume-level operations are disabled.

User response

Use other methods, such as 'Mount as iSCSI target', to restore data from the selected snapshot.

FBP7020E **The partitions in the selected disk are not formatted with a supported file system.**

Explanation

Recovery Agent successfully parsed the partition structure of the disk. However, none of the partitions use these supported file systems: FAT, NTFS, EXT2, EXT3, EXT4, or ReiserFS. As a result, volume-level operations, such as 'Instant Restore' and 'Mount as Virtual Volume', are not supported for this snapshot.

System action

Volume-level operations are disabled.

User response

Use other methods, such as 'Mount as iSCSI target', to restore data from the selected snapshot.

FBP7021E **Select the node access method.**

Explanation

Three node authentication methods are available to access snapshots on the Tivoli Storage Manager server: 'Asnodename' authenticates with a proxy node, 'Fromnode' authenticates with a node that contains limited access, and 'Direct' authenticates directly.

System action

The system waits for a user response.

User response

Specify the node authentication method to access the snapshots on the Tivoli Storage Manager Server. See information about these three methods in the product documentation.

FBP7022E	You must specify a Tivoli Storage Manager server name or IP.
-----------------	---

Explanation

Recovery Agent requires this information to access the virtual machines backup snapshots on the Tivoli Storage Manager Server.

System action

The system waits for a user response.

User response

Specify the host name or IP address of the Tivoli Storage Manager Server where the backup snapshots are located.

FBP7023E	You must specify a valid Tivoli Storage Manager Server port.
-----------------	---

Explanation

Recovery Agent requires this information to access the virtual machines backup snapshots on the Tivoli Storage Manager Server.

System action

The system waits for a user response.

User response

Specify the port number that is used by the Tivoli Storage Manager Server where the backup snapshots are located.

FBP7024E	You must specify an authentication node.
-----------------	---

Explanation

No authentication node was specified. Specify the Tivoli Storage Manager node that owns the snapshots.

System action

The system waits for a user response.

User response

Specify the Tivoli Storage Manager node that owns the snapshots. See information about the node authentication methods in the product documentation.

FBP7025E	You must specify a target node.
-----------------	--

Explanation

Recovery Agent provides three node authentication methods. When 'Fromnode' or 'Asnodename' are selected, a target node must be specified. The target node is the Tivoli Storage Manager node where the virtual machine backups are located.

System action

The system waits for a user response.

User response

Specify the target node where the virtual machine backups are located. See information about the node authentication methods in the product documentation.

FBP7026E	You must specify a node password.
-----------------	--

Explanation

No password was specified for the authentication node.

System action

The system waits for a user response.

User response

Enter the password of the Tivoli Storage Manager node that owns the virtual machine snapshots.

FBP7027E	The Recovery Agent 'Driver timeout' option requires a value of 60 through 300.
-----------------	---

Explanation

The Recovery Agent 'Driver timeout' option specifies the amount of time (in seconds) to process data requests from the file system driver. When processing does not complete within the specified time, the request is canceled and an error is returned to the file system driver.

System action

The system waits for a user response.

User response

Specify a valid 'Driver timeout' value of 60 through 300. See information about setting this option in the product documentation.

FBP7028E	The Recovery Agent 'Write cache size' option requires a value of 1 through <i>upper limit</i>
-----------------	--

Explanation

During Linux instant restore and mount operations, the Recovery Agent on the Windows backup proxy host saves data changes in the write cache folder. The maximum cache size is 90% of the available space for the selected folder.

System action

The system waits for a user response.

User response

Specify a valid Virtual Volume write 'Cache size' value. See information about setting this option in the product documentation.

FBP7029E	The Recovery Agent CLI mount command is missing one or more required parameters.
-----------------	---

Explanation

The Recovery Agent CLI mount command cannot complete without all required parameters.

System action

The operation is canceled.

User response

Issue the 'RecoveryAgentShell.exe -h mount' (Windows) or 'RecoveryAgentShell -h mount dump' (Linux) command to view the required parameters. Then, issue the mount command again with all required parameters.

FBP7030E	Repository '<i>repository name</i>' was not found
-----------------	--

Explanation

The Tivoli Storage Manager Server connection specified in the Recovery Agent '-rep' tag was not found.

System action

The operation is canceled.

User response

Correct the Recovery Agent CLI command '-rep' tag to identify the Tivoli Storage Manager Server where the backup snapshots are located.

FBP7031E	A valid iSCSI target name must be specified. Valid iSCSI names consist of the following items: lower-case English characters, digits, '.', ':', and '-'.
-----------------	---

Explanation

When a snapshot is exposed as an iSCSI target, a valid iSCSI target name must be entered.

System action

The operation is canceled.

User response

Specify a valid iSCSI target name.

FBP7032E	Mount target '<i>mount target</i>' is not valid.
-----------------	---

Explanation

The specified mount path is not a valid folder path.

System action

The operation is canceled.

User response

Specify a valid folder path.

FBP7033E	Reparsing point '<i>reparsing point</i>' was not found and cannot be created
-----------------	---

Explanation

The specified mount target path was not found and cannot be created.

System action

The operation is canceled.

User response

Specify a valid folder path.

FBP7035E Failed to load partition '*partition*'**Explanation**

The Recovery Agent was unable to retrieve the partition list from the disk snapshot.

System action

The operation is canceled.

User response

Check the Recovery Agent logs for information about why the partition list was not retrieved.

FBP7036E Incorrect partition number '*partition number*'**Explanation**

The specified partition was not found on the disk snapshot.

System action

The operation is canceled.

User response

Specify a valid partition number.

FBP7037E '*Cache size*' must be at least 1GB.**Explanation**

Mount and instant restore sessions that run on a Linux machine use the Recovery Agent Cache for write operations.

System action

The system waits for a user response.

User response

Specify a valid size for the 'Write Cache' size field. '*Cache size*' must be at least 1GB.

FBP7038E The value of the Recovery Agent write '*Cache size*' option must not exceed *max size in GBGB***Explanation**

During Linux instant restore and mount operations, the Recovery Agent on the Windows backup proxy host saves data changes in the write cache folder. The

maximum cache size is 90% of the available space for the selected folder.

System action

Linux mount and instant restore sessions might fail when the cache size reaches its limit.

User response

Do not start a new Recovery Agent mount or instant restore session on the Linux machine until the Virtual Volume write '*Cache size*' value decreases. See information about setting this option in the product documentation.

FBP8001E resume failed**Explanation**

The instant restore session is unable to resume. As a result, the instant restore session pauses. The restored volume is inaccessible while the session is paused.

System action

The instant restore session pauses.

User response

Check the Recovery Agent logs for information about why the resume failed. After the issue is resolved, resume the session.

FBP8002E failed to mount because of too many mount points**Explanation**

The Recovery Agent supports a maximum of 128 simultaneously mounted snapshots. That maximum was exceeded.

System action

The mounting operation is canceled.

User response

Dismount at least one of the currently mounted snapshots before an attempt to mount another snapshot.

FBP8003E failed to dismount *mounted volume name*

Explanation

The Recovery Agent failed to dismount the mounted snapshot. This issue might be caused by a disconnection from the Tivoli Storage Manager server that owns the snapshots.

System action

The dismount operation is canceled.

User response

Check the Recovery Agent logs for information about why the dismount operation failed.

FBP8004E **failed to load *repository name***

Explanation

The Recovery Agent failed to expose the snapshots of the specified Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Verify that the correct Tivoli Storage Manager server and node that own the snapshots are specified.

FBP8007E **Virtual Volume Driver not enabled**

Explanation

Mount operations require a working Recovery Agent Virtual Volume Driver.

System action

The mount operations are canceled.

User response

Check the Recovery Agent logs for information about why the Virtual Volume Driver is not enabled.

FBP8008E **snapshot not found.**

Explanation

The selected snapshot was not found on the Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Click Refresh in the Recovery Agent GUI to load the current snapshots on the Tivoli Storage Manager server.

FBP8009E **already mounted**

Explanation

The selected snapshot was already mounted to the requested target.

System action

None.

User response

The mounted snapshot is available for use.

FBP8012E **target is a network-mapped drive**

Explanation

Mounting snapshots to a network-mapped drive is not supported.

System action

The mount operation is canceled.

User response

Specify a target drive that is not on a network-mapped drive.

FBP8015E **volume letter is in use. Select another**

Explanation

The selected drive letter for the mount operation is in use.

System action

The mount operation is canceled.

User response

Select a target drive letter that is not in use.

FBP8016E **failed to mount**

Explanation

The Recovery Agent failed to complete the mount operation.

System action

The mount operation is canceled.

User response

Check the Recovery Agent logs for information about why the mount operation failed to complete.

FBP8019E **failed to stop**

Explanation

The current request to abort the selected instant restore session failed because the Recovery Agent could not locate the selected session.

System action

None.

User response

The abort request for the instant restore session was already done.

FBP8020E **failed to finalize the session**

Explanation

The product encountered an internal error when it attempted to finalize the instant restore session.

System action

The operation is canceled.

User response

Check the Recovery Agent logs for information about why the session did not finalize.

FBP8023E **target partition is too small**

Explanation

The volume size of the destination location must be equal to, or greater than, the size of the original volume on the snapshot to be restored.

System action

The instant restore operation is canceled.

User response

Specify a target volume with a size equal to, or greater than, the source volume on the snapshot to be restored.

FBP8024E **load repository *repository name* failed: *error message***

Explanation

The Recovery Agent failed to expose the snapshots of the specified Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Verify that the correct Tivoli Storage Manager server and Node that own the snapshots are specified.

FBP8025E **repository inaccessible.
Dismounting volume**

Explanation

The Recovery Agent is unable to retrieve data from the Tivoli Storage Manager server. As a result, the currently mounted snapshots are dismounted.

System action

Mounted snapshots are dismounted.

User response

Check the Recovery Agent logs for information about why the repository is inaccessible.

FBP8026E **'*path to repository*' inaccessible or not a repository**

Explanation

The Recovery Agent failed to expose the snapshots of the specified Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Verify that the correct Tivoli Storage Manager server and Node that own the snapshots are specified.

FBP8027E **failed to open repository**

Explanation

The Recovery Agent failed to expose the snapshots of the specified Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Verify that the correct Tivoli Storage Manager server and Node that own the snapshots are specified.

FBP8029E session stopped by user

Explanation

The user requested to abort the instant restore session. Aborting the instant restore sessions causes all data that was written to the restored volume to be lost.

System action

The instant restore session ends.

User response

The restored volume is shown as unformatted and requires reformatting.

FBP8031E Exclusive access to the mounted snapshot was not obtained on the Tivoli Storage Manager server.

Explanation

An exclusive access to the snapshot data on the Tivoli Storage Manager server could not be obtained. As a result, the version being restored could expire, leading to inability to complete the restore. Failure to obtain exclusive access is often the result of the snapshot data residing on a target replication server.

System action

The mount operation is canceled.

User response

If expiration protection is enabled, check the status of the target Tivoli Storage Manager server. If the target server is the replication server in failover mode, or if you verified no snapshots are running on the primary server, disable expiration protection. Then, try the operation again. If expiration protection was disabled when this error occurred, visit the IBM Support Portal for additional information at <http://www.ibm.com/support/entry/portal/>.

FBP8032E failed to unmount volume

Explanation

The target volume for the restore operation is in use. As a result, the instant restore operation did not start. Restoring a volume to a viewable storage volume involves overwriting data on that existing storage volume. After the restore session begins, the data on the existing volume is permanently erased.

System action

The instant restore session is canceled.

User response

Close any application (such as Windows Explorer or a command prompt) that might be accessing this volume. Then, try the operation again.

FBP8033E failed to restore blocks

Explanation

The instant restore operation is either unable to retrieve data from the Tivoli Storage Manager server or unable to write data to the target volume.

System action

The instant restore session ends.

User response

Check the Recovery Agent logs for information about why the blocks failed to restore. The restored volumes display as unformatted and require reformatting.

FBP8034E failed to mount volume

Explanation

The Recovery Agent failed to mount the target volume and start the instant restore operation.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information regarding why the target volume could not be mounted.

FBP8036E Failed to finalize

Explanation

The Recovery Agent failed to finalize the instant restore session.

System action

The instant restore session is canceled.

User response

Check the Recovery Agent logs for information about the cause of the problem. The restored volumes might display as unformatted and require reformatting.

FBP8037E **initialization failed. See logs for the reason**

Explanation

The Recovery Agent failed to initialize the instant restore operation.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information about the cause of the problem.

FBP8041E **cannot restore to a dynamic disk**

Explanation

Instant restore to a dynamic volume is not supported.

System action

The instant restore operation is canceled.

User response

Select a basic volume as the instant restore target. Then, try the operation again.

FBP8042E **cannot restore to clustered disk**

Explanation

Instant restore of a volume in a clustered environment is supported. However, the user canceled the restore to a volume in a clustered environment.

System action

The instant restore operation is canceled.

User response

Select a different volume as the instant restore target. Then, try the operation again.

FBP8043E **failed to create bitmap**

Explanation

The Recovery Agent failed to create the required internal data structure for the instant restore operation.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information about why the data structure was not created.

FBP8044E **failed to scramble first block**

Explanation

The Recovery Agent failed to overwrite the first sector of the disk.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information about why the overwrite to disk failed.

FBP8045E **failed to notify driver**

Explanation

The Recovery Agent failed to notify the kernel driver regarding the start of the instant restore session.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information about why the driver was not notified.

FBP8046E **failed to unscramble first block**

Explanation

The Recovery Agent failed to overwrite the first sector of the disk.

System action

The instant restore session ends.

User response

Check the Recovery Agent logs for information about why the overwrite to disk failed. The restored volumes might display as unformatted and require reformatting.

FBP8047E **cannot restore to a FAT volume.
Format target volume as NTFS**

Explanation

Instant restore to a volume formatted as an FAT32 file system is not supported.

System action

The instant restore operation is canceled.

User response

Format the volume as an NTFS file system. Then, try the operation again.

FBP8048E **session not responding**

Explanation

The instant restore session did not respond to the abort request within 5 minutes. As a result, the instant restore session was forcibly stopped.

System action

The instant restore session ends.

User response

Check the Recovery Agent logs for information about why the Recovery Agent did not respond to the abort request. The restored volumes might display as unformatted and require reformatting.

FBP8050E **failed to create first block file**

Explanation

The Recovery Agent failed to create a memory-mapped file for the instant restore session.

System action

The instant restore operation is canceled.

User response

Check the Recovery Agent logs for information regarding why the memory-mapped file was not created.

FBP8051E **cannot restore to disk with
signature '0'**

Explanation

Instant restore is not supported for disks without an MBR disk signature.

System action

The instant restore operation is canceled.

User response

Select an instant restore target volume on an MBR disk that contains a disk signature.

FBP8052E **Recovery Agent is currently
initializing.**

Explanation

Mount and instant restore operations cannot proceed when the Recovery Agent is initializing.

System action

The mount or instant restore operation is canceled.

User response

Wait until the Recovery Agent completes initializing. Then, try the operation again.

FBP8053E **failed to read data from server**

Explanation

The Recovery Agent failed to expose the snapshots of the specified Tivoli Storage Manager server.

System action

The operation is canceled.

User response

Verify that the correct Tivoli Storage Manager server and Node that own the snapshots are specified. Check the Recovery Agent logs for more information.

FBP9000E **failed connecting to a kernel driver**

Explanation

Instant restore operations require a working Recovery Agent kernel driver.

System action

The instant restore operations are canceled.

User response

Check the Recovery Agent logs for information about why the Recovery Agent failed to connect to the driver.

FBP9001E	Incompatible Virtual Volume Driver (FBVV) Version , expecting <i>expected major version.expected minor version</i> , installed <i>installed major version.installed minor version</i>
-----------------	--

Explanation

The kernel driver version does not match the Recovery Agent version. A valid driver is required for the Recovery Agent to work properly.

System action

The Recovery Agent application closes.

User response

The Recovery Agent was not installed correctly. Follow the Recovery Agent installation instructions in the product documentation.

FBP9002E	Cannot initialize Windows Sockets.
-----------------	---

Explanation

The Recovery Agent failed to initialize the Windows Sockets DLL file.

System action

The Recovery Agent application closes.

User response

Check the Windows events logs for errors related to this issue. Also check the Recovery Agent logs for information about why the Windows Sockets DLL file failed to initialize.

FBP9003E	Cannot obtain the application data directory path
-----------------	--

Explanation

The Recovery Agent was unable to retrieve the application data directory path from the operating system.

System action

The Recovery Agent application closes.

User response

Check the Windows events log for errors related to this issue. Also check the Recovery Agent logs for information about why the Recovery Agent was unable to obtain the application data directory path.

FBP9004E	Cannot create directory <i>directory name</i>
-----------------	--

Explanation

The Recovery Agent failed to create the application data directory.

System action

The Recovery Agent application closes.

User response

Check the Windows events logs for any errors. Check the Recovery Agent logs for information about why the Recovery Agent failed to create the application data directory.

FBP9005E	failed to initialize <i>module name</i> module
-----------------	---

Explanation

The Recovery Agent failed to initialize.

System action

The Recovery Agent application closes.

User response

Check the Recovery Agent logs for information about why the Recovery Agent failed to initialize.

FBP9006E	another instance of Recovery Agent is already running
-----------------	--

Explanation

Only one active Recovery Agent instance is supported. This issue is encountered when multiple users are logged on to the system and attempting to run the Recovery Agent.

System action

The Recovery Agent is not started.

User response

Either close the current Recovery Agent instance or run the Recovery Agent on a different system.

FBP9007E	unable to install the Recovery Agent
-----------------	---

Explanation

The Recovery Agent failed to install. A valid installation is required for the Recovery Agent to function properly.

System action

The Recovery Agent is not started.

User response

Follow the Recovery Agent installation instructions in the product documentation.

FBP9008E	Cannot get folder name for AFS.dll
-----------------	---

FBP9009E	Registration of Virtual Volume Driver failed .Check the file <i>file name</i>\\installFBVV.log for more information Do you want to retry registering the Virtual Volume Driver?
-----------------	--

Explanation

The Virtual Volume Driver must be registered correctly in order for the Recovery Agent to function correctly.

System action

The system waits for a user response.

User response

>Check the Recovery Agent logs for information about why the Recovery Agent failed to register the driver. Click 'Retry' to make another attempt to register the driver or click 'Cancel' to end the operation.

FBP9010E	Write Cache is full.
-----------------	-----------------------------

Explanation

Mount and instant restore sessions that run on a Linux machine use the Recovery Agent Virtual Volume write cache for write operations. Linux mount and instant restore sessions might fail when the cache size reaches its limit.

System action

Write operations to the Linux instant restore and virtual volumes might be lost.

User response

Unmount some of the mounted volumes on the Linux machine to make available space in the write cache. Instant restore volumes on the Linux machine might display as unformatted. When the cache is full, all data that is written to the Linux instant restore volumes is lost.

Appendix E. Accessibility features for the IBM Storage Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Storage Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Storage Protect product family 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 accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](http://www.ibm.com/able) (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119
Armonk, NY 10504-1785
US*

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 wish 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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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 document 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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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 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_.

Trademarks

IBM, the IBM logo, and 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 www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Storage Protect family of products.

See the [IBM Storage Protect glossary](#).

Index

Special Characters

.vmx file
attributes [176](#)

A

accessibility features [245](#)
Active Directory
 verify replication [158](#)
Active Directory domain controllers [93](#)
application protection
 Active Directory domain controllers [93](#)
 configure backup policy [36](#)
 USN Rollback [93](#)
automated client failover
 overview [13](#)

B

backing up virtual machine data
 with Data Protection for VMware [129](#)
backing up VM data
 with IBM Storage Protect Snapshot for VMware [132](#)
backup
 backing up with one data mover [141](#), [160](#), [161](#)
 migrated virtual machine [133](#)
 organization vDC [134](#)
 specifying domain-level parameters [139](#)
 specifying objects [149](#)
 task [133](#), [139](#), [141](#), [149](#), [160](#), [161](#)
 templates [12](#)
 vApp [12](#)
 vmcli command [97](#)
backup management
 configure backup policy [28](#)
backup policy
 change retention policy [32](#)
 change snapshot attempts [35](#)
 configure [28](#)
 enable application protection [36](#)
 exclude virtual machines [31](#)
 include virtual machines [31](#)
 include VM disks [34](#)
 set data consistency [35](#)
 set data mover [33](#)
 set disk protection [34](#)
 set schedule [30](#)

C

change retention policy
 configure backup policy [32](#)
commands
 data mover [95](#)
configure IBM Storage Protect vSphere Client plug-in

 configure IBM Storage Protect vSphere Client plug-in (*continued*)
 create data protection tags [25](#)
 set data mover node [22](#)
 tagging support [22](#)
create tags
 configure IBM Storage Protect vSphere Client plug-in [25](#)

D

data consistency
 configure backup policy [35](#)
data mover
 commands [95](#)
 options [95](#)
 reference [95](#)
data movers
 edit in schedule [38](#), [44](#), [46](#)
data protection
 configure [28](#)
 general help [64](#)
Data Protection for VMware
 using Data Protection for Microsoft Exchange Server [64](#)
disability [245](#)
disks
 control [12](#)
domain controller
 verify replication [158](#)

E

enable tracing
 troubleshooting the IBM Storage Protect vSphere Client
 plug-in [179](#), [180](#)
errors [163](#)
exclude virtual machines
 configure backup policy [31](#)

F

failover
 client [13](#)
FBP1001I [222](#)
FBP1003I [223](#)
FBP1005I [223](#)
FBP1008I [223](#)
FBP1009I [223](#)
FBP1011I [223](#)
FBP1013I [224](#)
FBP1014I [224](#)
FBP1100I [224](#)
FBP1101I [224](#)
FBP1102I [224](#)
FBP1103I [224](#)
FBP1104I [224](#)
FBP1300I [225](#)
FBP1301W [225](#)

[FBP1302E](#) [225](#)
[FBP1303E](#) [225](#)
[FBP1304E](#) [225](#)
[FBP1305I](#) [225](#)
[FBP5003W](#) [225](#)
[FBP5005W](#) [226](#)
[FBP5007W](#) [226](#)
[FBP5008W](#) [226](#)
[FBP5010W](#) [226](#)
[FBP5011W](#) [227](#)
[FBP5012W](#) [227](#)
[FBP5013W](#) [227](#)
[FBP5015W](#) [227](#)
[FBP5017W](#) [227](#)
[FBP5018W](#) [228](#)
[FBP5020W](#) [228](#)
[FBP5021W](#) [228](#)
[FBP5023W](#) [228](#)
[FBP5025W](#) [229](#)
[FBP5026W](#) [229](#)
[FBP5028W](#) [229](#)
[FBP5029W](#) [229](#)
[FBP5030W](#) [229](#)
[FBP5031W](#) [230](#)
[FBP5032W](#) [230](#)
[FBP5033W](#) [230](#)
[FBP5034W](#) [230](#)
[FBP5035W](#) [230](#)
[FBP7003E](#) [231](#)
[FBP7004E](#) [231](#)
[FBP7005E](#) [231](#)
[FBP7006E](#) [231](#)
[FBP7007E](#) [231](#)
[FBP7008E](#) [231](#)
[FBP7009E](#) [232](#)
[FBP7012E](#) [232](#)
[FBP7013E](#) [232](#)
[FBP7014E](#) [232](#)
[FBP7015E](#) [232](#)
[FBP7016E](#) [232](#)
[FBP7017E](#) [233](#)
[FBP7018E](#) [233](#)
[FBP7019E](#) [233](#)
[FBP7020E](#) [233](#)
[FBP7021E](#) [233](#)
[FBP7022E](#) [234](#)
[FBP7023E](#) [234](#)
[FBP7024E](#) [234](#)
[FBP7025E](#) [234](#)
[FBP7026E](#) [234](#)
[FBP7027E](#) [234](#)
[FBP7028E](#) [235](#)
[FBP7029E](#) [235](#)
[FBP7030E](#) [235](#)
[FBP7031E](#) [235](#)
[FBP7032E](#) [235](#)
[FBP7033E](#) [235](#)
[FBP7035E](#) [236](#)
[FBP7036E](#) [236](#)
[FBP7037E](#) [236](#)
[FBP7038E](#) [236](#)
[FBP8001E](#) [236](#)
[FBP8002E](#) [236](#)
[FBP8003E](#) [236](#)

[FBP8004E](#) [237](#)
[FBP8007E](#) [237](#)
[FBP8008E](#) [237](#)
[FBP8009E](#) [237](#)
[FBP8012E](#) [237](#)
[FBP8015E](#) [237](#)
[FBP8016E](#) [237](#)
[FBP8019E](#) [238](#)
[FBP8020E](#) [238](#)
[FBP8023E](#) [238](#)
[FBP8024E](#) [238](#)
[FBP8025E](#) [238](#)
[FBP8026E](#) [238](#)
[FBP8027E](#) [238](#)
[FBP8029E](#) [239](#)
[FBP8031E](#) [239](#)
[FBP8032E](#) [239](#)
[FBP8033E](#) [239](#)
[FBP8034E](#) [239](#)
[FBP8036E](#) [239](#)
[FBP8037E](#) [240](#)
[FBP8041E](#) [240](#)
[FBP8042E](#) [240](#)
[FBP8043E](#) [240](#)
[FBP8044E](#) [240](#)
[FBP8045E](#) [240](#)
[FBP8046E](#) [240](#)
[FBP8047E](#) [241](#)
[FBP8048E](#) [241](#)
[FBP8050E](#) [241](#)
[FBP8051E](#) [241](#)
[FBP8052E](#) [241](#)
[FBP8053E](#) [241](#)
[FBP9000E](#) [241](#)
[FBP9001E](#) [242](#)
[FBP9002E](#) [242](#)
[FBP9003E](#) [242](#)
[FBP9004E](#) [242](#)
[FBP9005E](#) [242](#)
[FBP9006E](#) [242](#)
[FBP9007E](#) [243](#)
[FBP9008E](#) [243](#)
[FBP9009E](#) [243](#)
[FBP9010E](#) [243](#)

file restore

[configuring tracing](#) [169](#)

[description](#) [51](#)

[logging in](#) [54](#)

[prerequisites](#) [51](#)

[procedure](#) [54](#)

[solutions](#)

[unique issues](#) [169](#)

[troubleshooting](#)

[diagnostic procedure](#) [168](#)

files

[restore overview](#) [186](#)

full VM instant restore

[environment requirements](#) [14](#)

[scenarios](#) [154](#)

[validation scenarios](#) [157](#)

G

[get_password_info](#)

get_password_info (continued)
vmcli command [112](#)

I

IBM Documentation [vii](#)
IBM Storage Protect vSphere Client plug-in
about [19](#)
available features [21](#)
canceling backups [42](#)
connecting to the Data Protection for VMware vSphere GUI [22](#)
dismounting a virtual machine [49](#)
getting started [19](#)
restoring virtual machines [47](#)
schedules. managing [38](#), [44](#), [46](#)
schedules. viewing run history [39](#)
setting at-risk policy for virtual machines [43](#)
starting on-demand backups [40](#)
troubleshooting
messages [181](#)
viewing backup history for virtual machines [43](#)
viewing backup operations for virtual machines [42](#)
include virtual machines
configure backup policy [31](#)
inquire_config
vmcli command [104](#)
inquire_detail
vmcli command [106](#)
instant restore
overview [186](#)
task (Windows) [188](#)

K

keyboard [245](#)

L

LAN environment [185](#)

M

mailbox history information
updating in Microsoft Exchange Server backups [63](#)
manage
schedules [38](#), [44](#), [46](#)
manage data protection [28](#)
messages
Data Protection for VMware vSphere GUI [191](#)
recovery agent [213](#)
Microsoft Exchange Server backups
updating mailbox history [63](#)
mounting a disk [151](#)
mounting snapshots [185](#)

N

New in Data Protection for VMware 8.1.27 [ix](#)

O

operating systems
Windows [185](#)
options
data mover [95](#)
organization vDC
backup [134](#)
out-of-space errors [14](#)

P

platform services controller connection
troubleshooting the IBM Storage Protect vSphere Client
plug-in [178](#)
problem determination [163](#)
publications [vii](#)

R

replication
verify Active Directory [158](#)
verify domain controller [158](#)
restore
configuring tracing [169](#)
file
solutions [169](#)
troubleshooting [168](#)
file restore description [51](#)
instant
full VM requirements [14](#)
logging in [54](#)
prerequisites [51](#)
procedure [54](#)
templates [12](#)
vApp [12](#)
vmcli command [99](#)
vSphere scenario [153](#)
restoring data
Exchange Server 2010 [73](#)
Exchange Server 2013 [73](#)
Mailbox Restore Browser [73](#)

S

schedules
edit data movers [38](#), [44](#), [46](#)
managing [38](#), [44](#), [46](#)
scheduling a backup
with Data Protection for VMware [129](#)
with IBM Storage Protect Snapshot for VMware [132](#)
set data mover
configure backup policy [33](#)
set disk protection
configure backup policy [34](#)
set schedule
configure backup policy [30](#)
set tag data mover
configure IBM Storage Protect vSphere Client plug-in [22](#)
set_domian
vmcli command [108](#)
set_option
vmcli command [109](#)

- set_password
 - vmcli command [110](#)
- snapshot attempts
 - configure backup policy [35](#)
- snapshots
 - mounting [185](#)
- SSL [177](#)
- start_guest_scan
 - vmcli command [113](#)

T

- tagging support
 - configure IBM Storage Protect vSphere Client plug-in [22](#)
- templates [12](#)
- tracing
 - file restore [169](#)
- troubleshooting
 - digital certificates [177](#)
 - file restore
 - diagnostic procedure [168](#)
 - unique issues [169](#)
 - IBM Storage Protect vSphere Client plug-in problems [178](#)
- Troubleshooting file restore operations [170](#), [172](#), [176](#)
- troubleshooting the IBM Storage Protect vSphere Client plug-in
 - enable tracing [179](#), [180](#)
 - messages [181](#)
 - platform services controller connection [178](#)

U

- updating mailbox history information [63](#)
- USN Rollback [93](#)

V

- vApp [12](#)
- Virtual Volumes (VVOL), about [10](#)
- VM backup [131](#)
- vmcli command
 - backup [97](#)
 - get_password_info [112](#)
 - inquire_config [104](#)
 - inquire_detail [106](#)
 - restore [99](#)
 - set_domain [108](#)
 - set_option [109](#)
 - set_password [110](#)
 - start_guest_scan [113](#)
- vmdatastorethreshold
 - usage [14](#)
- volumes
 - restore overview [186](#)
- VSS backup Data Protection for VMware
 - with Data Protection for Microsoft Exchange Server [64](#)



Product Number: 5725-X00