

IBM Storage Protect for Virtual Environments
8.1.27

*Data Protection for Microsoft Hyper-V
Installation and User's Guide*



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 265](#).

This edition applies to version 8, release 1, modification 27 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 for version 8.1.27.....	ix
Chapter 1. Protection for Microsoft Hyper-V virtual machines.....	1
Back up Hyper-V virtual machines.....	1
VM backups with Volume Shadow Copy Service (VSS).....	1
VM backups with resilient change tracking (RCT).....	2
Restore Hyper-V virtual machines.....	3
User interfaces for Hyper-V operations.....	4
How IBM Storage Protect nodes are used.....	6
Policy management at the virtual machine level.....	8
Incremental forever backup strategy.....	8
Snapshot management with Windows PowerShell.....	8
Automated client failover for Data Protection for Hyper-V.....	9
Limitations on Hyper-V backup operations	10
Documentation resources.....	11
Chapter 2. Installing and upgrading Data Protection for Microsoft Hyper-V.....	15
Planning to install Data Protection for Microsoft Hyper-V.....	15
Features that are installed.....	15
System requirements.....	15
Required communication ports.....	16
Upgrading Data Protection for Microsoft Hyper-V.....	17
Compatibility with different versions.....	17
Renaming nodes on the IBM Storage Protect server.....	18
Customizing node names.....	20
Upgrade considerations for RCT backups.....	22
Migrating from VSS backups to RCT backups.....	22
Upgrading by using the installation wizard.....	23
Installing Data Protection for Microsoft Hyper-V components.....	23
Download and extract the installation package.....	23
Installing with the wizard.....	24
Uninstalling Data Protection for Microsoft Hyper-V.....	29
Installing in silent mode.....	30
Uninstalling in silent mode.....	31
Installing on Windows Server Core systems.....	32
Uninstalling on Windows Server Core systems.....	33
Installing the Linux mount proxy.....	33
Uninstalling the Linux mount proxy feature.....	37
Removing the file restore feature.....	38
Chapter 3. Configuring Data Protection for Microsoft Hyper-V.....	39
Configuring Data Protection for Microsoft Hyper-V with the wizard.....	39
Configuring security settings for Data Protection for Microsoft Hyper-V.....	42
Configuring security settings to connect to IBM Storage Protect server version 8.1.1 or earlier or version 7.1.7 or earlier.....	43
Enabling the environment for file restore operations.....	44

Configuring the Linux mount proxy for file restore operations.....	45
Checking the configuration of the Linux file restore environment.....	47
Modifying options for file restore operations.....	48
Options for file restore operations.....	48
Configuring Data Protection for Microsoft Hyper-V log activity.....	49
Data Protection for Microsoft Hyper-V log activity options.....	50
Configuring the IBM Storage Protect recovery agent GUI	50
Enabling secure communication from the recovery agent to the IBM Storage Protect server.....	54
Manually configuring an iSCSI device.....	57
Advanced configuration.....	58
Configuring non-default port numbers for Data Protection for Microsoft Hyper-V operations.....	58
Tuning scheduled VM backups for Windows Server 2012 clusters.....	59

Chapter 4. Managing data with the Data Protection for Microsoft Hyper-V

Management Console.....	61
Starting the management console.....	61
Navigating the management console.....	62
Navigation pane.....	62
Results pane.....	63
Actions pane.....	71
Verifying the configuration of Data Protection for Microsoft Hyper-V.....	71
Managing backup schedules for a host or cluster machine.....	72
Setting the at-risk policy for a virtual machine.....	74
Viewing the schedule history for a Hyper-V host or cluster.....	75
Viewing the backup status and backup history of a virtual machine.....	75
Running an ad hoc backup of a virtual machine.....	77
Restoring a virtual machine.....	78
Best practices for Data Protection for Microsoft Hyper-V.....	79

Chapter 5. Getting started with file restore operations.....81

File restore tasks.....	81
File restore prerequisites.....	82
Logging in to restore files.....	84
Restoring files from a virtual machine backup.....	84

Chapter 6. Protecting in-guest applications..... 87

Protecting Microsoft Exchange Server data in Hyper-V environments	87
Installing and configuring software.....	87
Managing backup operations.....	96
Restoring data.....	99
IBM Storage Protect file space information.....	102
Protecting Microsoft SQL Server data in Hyper-V environments	103
Installing and configuring software.....	103
Managing backup operations.....	112
Restoring data.....	115
Sample script for validating full virtual machine backups.....	119
IBM Storage Protect file space information.....	120
Troubleshooting application protection of guest virtual machines.....	121
Troubleshooting VSS backup and restore operations on guest virtual machines.....	122

Chapter 7. Running instant access operations.....125

Verifying the integrity of a VM backup from the command line.....	125
Freeing instant access resources by using the command line.....	128
Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console.....	130
Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console.....	131

Restrictions on instant access VMs after a Hyper-V host is restarted.....	132
Chapter 8. Protecting virtual machines by using Windows PowerShell cmdlets....	133
Preparing to use PowerShell cmdlets.....	133
PowerShell cmdlets for Data Protection for Microsoft Hyper-V.....	135
Tasks that you can run with PowerShell cmdlets.....	138
Chapter 9. Command reference.....	145
Reading syntax diagrams.....	145
Backup VM	147
Expire	154
Query VM	154
Restore VM	158
Preview virtual machine restore operations.....	164
Chapter 10. Options reference.....	167
Dateformat.....	167
Detail.....	169
Domain.vmfull.....	169
Exclude.vmdisk.....	172
Inactive.....	174
Include.vm.....	174
Include.vmdisk.....	176
INCLUDE.VMSNAPSHOTATTEMPTS.....	177
INCLUDE.VMTSMVSS.....	179
Shadow copy considerations for restoring an application protection backup from the data mover.....	180
Mode.....	182
Mbobjrefreshtresh.....	183
Mbpctrefreshtresh.....	183
Noprompt.....	184
Numberformat.....	184
Pick.....	186
Pitdate.....	186
Pittime.....	187
Skipsystemexclude.....	187
Timeformat.....	188
Vmautostartvm.....	189
Vmbackdir.....	190
Vmctlmc.....	190
Vmiscsiserveraddress.....	191
Vmmaxbackupsessions.....	192
Vmmaxparallel.....	194
Vmmaxperssnapshot.....	195
Vmmaxrestoreparalleldisks.....	196
Vmmaxrestoreparallelvms.....	197
Vmmaxrestoresessions.....	202
Vmmaxsnapshotretry.....	203
Vmmaxvirtualdisks.....	204
Vmmc.....	206
Vmprocessvmwithphysdisks.....	206
Vmskipmaxvirtualdisks.....	207
Vmskipphysdisks.....	208
Vmrestoretype.....	209
Chapter 11. Mount and file restore	211
IBM Storage Protect recovery agent configurations.....	211

Snapshot mount overview.....	212
Mount guidelines.....	212
File restore overview.....	213
File restore guidelines.....	214
Restoring one or more files.....	214
Chapter 12. IBM Storage Protect recovery agent commands.....	217
Mount.....	217
Set_connection.....	221
Help.....	221
Recovery agent command-line interface return codes.....	222
Chapter 13. Optimizing performance.....	225
Optimizing backup operations.....	225
Backing up multiple virtual machines in parallel (optimized backup).....	227
Optimizing restore operations.....	227
Restoring multiple virtual machines in parallel (optimized restore).....	229
Appendix A. Troubleshooting.....	231
Troubleshooting Data Protection for Microsoft Hyper-V operations.....	234
Trace options for Data Protection for Microsoft Hyper-V.....	235
Appendix B. Data Protection for Microsoft Hyper-V messages.....	237
Appendix C. Accessibility.....	263
Notices.....	265
Glossary.....	269
Index.....	271

About this publication

This publication provides overview, planning, and user instructions for IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

Who should read this publication

This publication is intended for administrators and users who are responsible for implementing a backup solution with IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V in one of the supported environments.

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

- Microsoft Windows Server 2019 with the Hyper-V role installed
- Microsoft Windows Server 2016 with the Hyper-V role installed
- Microsoft Windows Server 2012 or 2012 R2 with the Hyper-V role installed
- The IBM Storage Protect backup-archive client
- The IBM Storage Protect server

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 for version 8.1.27

IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V version 8.1.27 introduces updates.

New and changed information in this product documentation is indicated by a vertical bar (|) to the left of the change.

The following updates are available for this release:

IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V now supports the Windows Server 2025 operating system

For a list of new features and updates for the current and previous version 8.1 releases, see [Data Protection for Microsoft Hyper-V updates](#).

Chapter 1. Protection for Microsoft Hyper-V virtual machines

IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V is a licensed product that provides storage management services for virtual machines (VMs) in a Microsoft Hyper-V environment.

Data Protection for Microsoft Hyper-V works with the IBM Storage Protect backup-archive client to protect Hyper-V virtual machines on the following operating systems:

- Microsoft Windows Server 2012
- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2016
- Microsoft Windows Server 2019

Back up Hyper-V virtual machines

Data Protection for Microsoft Hyper-V creates an incremental-forever full or incremental-forever incremental backup of Hyper-V virtual machines (VMs). A consistent snapshot is taken of the VM, and the VM is backed up to the IBM Storage Protect server.

You can back up Hyper-V VMs that exist on a local disk, a storage area network (SAN) attached disk, or Cluster Shared Volume (CSV). For example, you can back up VMs that are stored on CSVs in a Hyper-V cluster environment or on Server Message Block (SMB) file shares that are on a remote system. You can back up any guest operating systems that are supported by the Hyper-V server on remote shares, regardless of whether IBM Storage Protect supports them directly.

The following backup types are supported for Microsoft Hyper-V VMs with virtual disks that use the VHDX disk format:

Incremental-forever full backup

Creates a backup of snapshot disk data to the IBM Storage Protect server.

Incremental-forever incremental backup

Creates a snapshot of the blocks that changed since the last incremental-forever full backup or incremental-forever incremental backup.

If you are running the Hyper-V host on the Windows Server 2012 or Windows Server 2012 R2 operating system, Microsoft Volume Shadow Copy Service (VSS) is used to create a consistent snapshot of the VM. Changes that occur in the VM between each backup are tracked in a snapshot differencing file.

If you are running the Hyper-V host on a Windows Server 2016 or later operating system, snapshots are created by using a Windows API, and resilient change tracking (RCT) is used to track changes in a VHDX disk between each backup operation.

Virtual machine backups with Volume Shadow Copy Service (VSS)

For Hyper-V backups on Windows Server 2012 and 2012 R2, Microsoft Volume Shadow Copy Service (VSS) is used to create consistent snapshots of virtual machines (VMs) during backup operations.

During an initial incremental-forever full backup operation, the client creates a snapshot of the virtual machine hard disk (VHDX) and sends the content to the IBM Storage Protect server. Changes that occur after the initial snapshot are stored in a snapshot differencing file (.avhdx). Subsequent incremental-forever incremental backup operations back up only the data that was changed since the last backup.

If you run an incremental-forever incremental backup before you create an incremental-forever full backup, the client will run an incremental-forever full backup.

How snapshots work with VSS backups

During each VM backup, a new snapshot differencing file (.avhdx) is created to track the changes to the VM that occur after the backup operation. This differencing snapshot is saved on the Hyper-V host to collect the writes for the next incremental backup.

In previous releases of Data Protection for Microsoft Hyper-V, a snapshot could contain only one VM. This behavior could cause scheduling contention during cluster backup operations because too many snapshots had to be taken. By using the `vmmaxpersnapshot` option that was introduced in Data Protection for Microsoft Hyper-V version 8.1.2, you can reduce the number of snapshots that are taken for a backup operation by grouping several VMs in a single snapshot. For more information, see [“Tuning scheduled VM backups for Windows Server 2012 and 2012 R2 clusters”](#) on page 59.

Virtual machine backups with resilient change tracking (RCT)

For Hyper-V backups on Microsoft Windows Server 2016 or later versions, the resilient change tracking (RCT) feature is used to back up virtual machines (VMs).

RCT is a feature that provides built-in change block tracking capability for Hyper-V VM disks. Data Protection for Microsoft Hyper-V uses RCT to track changes to a VM disk (VHDX) that occur in between backup operations. The changes are tracked at the data block level. Only blocks that have changed since the last backup operation are candidates for the next incremental-forever incremental backup.

Windows Server 2016 also provides the capability to create backup snapshots (also known as checkpoints) directly without using Microsoft Volume Shadow Copy Service (VSS) (although VSS is still used inside Windows guest VMs to quiesce the VMs for application-consistent backups).

You can group several VMs in a single snapshot. However, if a guest VM that hosts applications is enabled for application protection, the VM snapshot is taken individually. For more information about application protection, see [Chapter 6, “Protecting in-guest applications,”](#) on page 87.

VM backup operations with RCT require the Hyper-V VM to be Version 6.2 or later.

If your VM was created on the Windows Server 2012 R2 or earlier operating system, and then later moved to a Windows Server 2016 host server (or the host server was upgraded to Windows Server 2016), you must take the VM offline and upgrade the VM version before it can be backed up. You can upgrade the VM Version by using the Hyper-V Manager or the **Update-VMVersion** cmdlet.

Data Protection for Microsoft Hyper-V version 8.1.0 uses VSS to back up VMs in the Hyper-V environment on Windows Server 2016. Starting in version 8.1.2, all Hyper-V VM backup operations in the Windows Server 2016 or later environment uses RCT. If you are upgrading from version 8.1.0, because previous VSS backups do not have RCT change-tracking information, the first time you use Data Protection for Microsoft Hyper-V version 8.1.27 to back up your VMs on Windows Server 2016, an incremental-forever full backup is created.

After you backed up a VM by using RCT, you can no longer use Data Protection for Microsoft Hyper-V version 8.1.0 to run VSS backups on that VM.

How snapshots work with RCT backups

During an incremental-forever full backup operation for a VM, a snapshot is created of the VM disk and the snapshot contents are backed up to the IBM Storage Protect server. The snapshot is deleted automatically after the backup operation is completed.

During the next incremental-forever incremental backup, a new snapshot is created and verified against the RCT change-tracking information from the previous backup operation to determine the data that has changed. Only the changed blocks are backed up to the IBM Storage Protect server.

After the backup operation, the snapshot is merged with the VM by Hyper-V, and the snapshot differencing file (.avhdx) is deleted automatically. This process is unlike the VSS snapshot processing on Windows Server 2012 and 2012 R2 operating systems, in which the snapshot differencing file is retained on the VM to store incremental changes.

Any snapshot that you create manually or with another backup product do not affect the backup chain that is created by the RCT process. You can create snapshots manually or with a third-party backup product before or after a Data Protection for Microsoft Hyper-V RCT backup operation, and the next incremental backup operation by Data Protection for Microsoft Hyper-V will be based on the RCT change-tracking information from the previous backup operation.

Features that are available for RCT backups

Most Data Protection for Microsoft Hyper-V features that work on Windows Server 2012 and 2012 R2 also apply to Windows Server 2016.

However, snapshot operations are different between VSS and RCT backups. For more information, see [“How snapshots work with RCT backups”](#) on page 2.

Support for host failover with Cluster Shared Volumes (CSVs) is unchanged from version 8.1.0 and earlier, but running a VM backup during a rolling upgrade of a Hyper-V cluster operating system is not supported.

How to query RCT backups

You can use the **query vm** command to display information about a VM that was backed up to the IBM Storage Protect server. Use the **-detail** parameter with the **query vm** command to show detailed information about the backup operation. For more information, see [“Query VM”](#) on page 154.

You can also use the **backup vm -preview** command to display the VM disk locations that can be used for the **backup vm** command. For more information, see [“Backup VM”](#) on page 147.

Related concepts

[“Upgrading Data Protection for Microsoft Hyper-V”](#) on page 17

Review the tasks that you need to do when you upgrade to Data Protection for Microsoft Hyper-V Version 8.1.27 from a previous version.

[“Limitations on Hyper-V backup operations”](#) on page 10

Before you start a Hyper-V backup operation, review the limitations. Some limitations apply to all Hyper-V backup operations, while others apply only to Hyper-V backups on Windows Server 2012 or 2012 R2 or Windows Server 2016 environments.

[“Virtual machine backups with Volume Shadow Copy Service \(VSS\)”](#) on page 1

For Hyper-V backups on Windows Server 2012 and 2012 R2, Microsoft Volume Shadow Copy Service (VSS) is used to create consistent snapshots of virtual machines (VMs) during backup operations.

Related tasks

[“Migrating from VSS backups to RCT backups”](#) on page 22

To take advantage of the resilient change tracking (RCT) backup feature in Data Protection for Microsoft Hyper-V version 8.1.2 or later, migrate your virtual machine (VM) backup operations from the Microsoft Volume Shadow Copy Service (VSS) to RCT.

Related reference

[“Troubleshooting”](#) on page 231

Solutions to Data Protection for Microsoft Hyper-V issues are provided.

Restore Hyper-V virtual machines

You can restore Hyper-V virtual machines (VMs) by using several methods. You can restore an entire virtual machine, restore an entire virtual machine to an alternative location, or restore individual files from a virtual machine.

Full VM Restore

Restore an entire Hyper-V VM

Each Hyper-V VM backup is restored from the IBM Storage Protect server as a single entity. You can restore any guest operating systems that are hosted by the Hyper-V server regardless of whether the guest operating system is supported by IBM Storage Protect.

A Data Protection for Microsoft Hyper-V restore operation ensures that the same block on the production disk is restored only once. Older backup versions expire according to the IBM Storage Protect server management class policy that is associated with the virtual machine.

Restore an entire Hyper-V VM to an alternative location

You can restore a Hyper-V VM to an alternative VM name, to an alternative location on the Hyper-V host, or both. You can also restore a Hyper-V VM to a different Hyper-V host by using the Data Protection for Microsoft Hyper-V Management Console. If you must use the command line to restore a VM to a different host, you must run the restore operation from the Hyper-V host where the VM is being restored to.

Restore files with the file restore interface

Use the IBM Storage Protect file restore interface to restore one or more files with a web-based interface. File owners can search, locate, and restore files from a VM backup with minimal administrator assistance. Help desk personnel can also use the file restore interface to restore files on behalf of file owners.

For more information, see [Chapter 5, “Getting started with file restore operations,” on page 81.](#)

Restore files with the recovery agent

Use this restore method only if you want to run in-guest mount operations. Files are manually copied from a mounted virtual machine disk that is accessed through an Internet Small Computer Systems Interface (iSCSI) target or partition. This method requires the IBM Storage Protect recovery agent to be installed.

For more information, see [Chapter 11, “Mount and file restore,” on page 211.](#)

Tip: Previous versions of Data Protection for Microsoft Hyper-V delivered functionality for mounting a VM disk and restoring a file with the recovery agent. Although this function is still supported by the recovery agent, the IBM Storage Protect restore interface is the preferred method.

Instant access

A temporary VM is created for verification of the backup data, but the VM is not restored.

You can use either the data mover command-line interface, the Data Protection for Microsoft Hyper-V Management Console, or Windows PowerShell cmdlets to run instant access operations.

For more information, see the following topics:

- [“Verifying the integrity of a VM backup from the command line” on page 125](#)
- [“Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console” on page 130](#)
- [“Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143](#)

User interfaces for Hyper-V operations

You can use several user interfaces to complete Data Protection for Microsoft Hyper-V Hyper-V operations. The data mover must be installed on the Hyper-V host server or on each host in a cluster.

The following user interfaces are available for Data Protection for Microsoft Hyper-V operations:

Data Protection for Microsoft Hyper-V Management Console

A graphical user interface that you can use to perform daily backup management tasks, such as managing virtual machine (VM) backups, monitoring VM backups, running ad hoc backup and restore operations, and updating the configuration.

IBM Storage Protect file restore interface

A web-based interface that file owners or help desk personnel can use to restore one or more files from a VM backup with minimal administrator assistance. The administrator provides a URL for the file restore interface.

Data mover

A component, also known as the backup-archive client, that moves data to and from the IBM Storage Protect server during backup and restore operations.

The data mover includes a command-line interface (**dsmc** commands) that you can use for backup, query, restore, and other operations.

Data Protection for Microsoft Hyper-V cmdlets

Windows PowerShell cmdlets that help you automate Data Protection for Microsoft Hyper-V operations with PowerShell scripts.

IBM Storage Protect recovery agent

An agent that provides virtual mount and file restore capability.

The following figures are the high-level overviews of Data Protection for Microsoft Hyper-V in the Windows Server 2016 or later and the Windows Server 2012 environments.

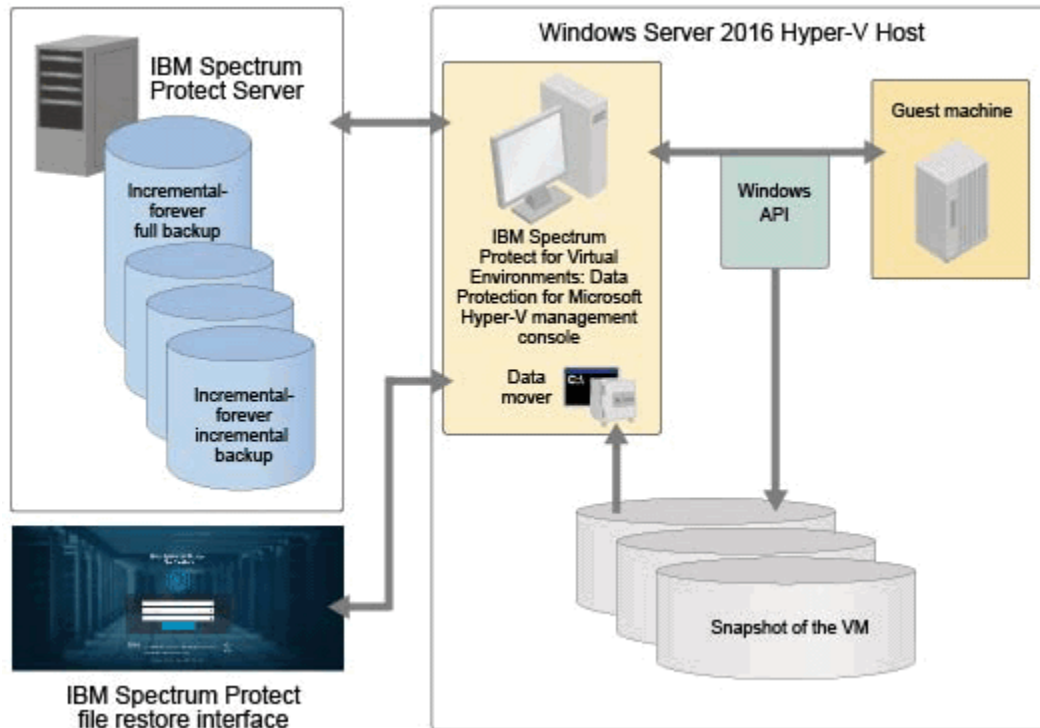


Figure 1. High-level overview of Data Protection for Microsoft Hyper-V in the Windows Server 2016 environment




Figure 2. High-level overview of Data Protection for Microsoft Hyper-V in the Windows Server 2012 environment

How IBM Storage Protect nodes are used in Data Protection for Microsoft Hyper-V

Data Protection for Microsoft Hyper-V 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 Microsoft Hyper-V, 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 Microsoft Hyper-V environment, the most basic node configuration consists of two nodes: the *data mover node* and the *target node*.

- The data mover node represents a specific data mover that "moves data" from one system to another. No data is stored under this node on the IBM Storage Protect server.
- The target is the node name under which VM data is stored on the IBM Storage Protect server.

In a cluster environment, the node configuration consists of a target node that is associated with the name of the cluster, and one data mover node for each host in the cluster.

For mount operations, a mount proxy node pair is required for each host system. 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 use the file restore interface to restore individual files, or use the recovery agent to 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 Hyper-V host node for each Windows or Linux system that serves as a proxy.

To simplify the configuration, the Data Protection for Microsoft Hyper-V configuration wizard automatically creates the various nodes that are required for backup, restore, and file restore operations. The configuration wizard also registers the nodes on the IBM Storage Protect server, creates the necessary

proxy relationships, creates the local options files, configures and starts the services for the data mover node on local Windows hosts.

The types of nodes that are created depend on your Hyper-V environment and whether you enabled the file restore feature. The node names that are created follow a specific naming convention that is based on the cluster or host name, and the node type. Custom node names cannot be used.

If you are upgrading from Data Protection for Microsoft Hyper-V version 8.1.2 or earlier and have nodes that are already defined on the IBM Storage Protect server, you must update the node names on the server. For more information, see [“Renaming nodes on the IBM Storage Protect server”](#) on page 18.

The following table contains a comparison of the different types of nodes in the Data Protection for Microsoft Hyper-V environment.

<i>Table 1. Types of nodes configured by the configuration wizard</i>		
Node type	Naming convention	Description
Target node	For a stand-alone host: <i>hostname_HV_TGT</i> For a cluster: <i>clustername_HV_TGT</i>	The node name where all VM backups are stored on the IBM Storage Protect server. For clusters, VMs are backed up to a single container on the IBM Storage Protect server under a single node name (cluster node), regardless of which host in the cluster is backing them up.
Data mover node	<i>hostname_HV_DM</i>	The node that backs up data to the target node on the IBM Storage Protect server. No data is stored under the data mover node. For clusters, a data mover node is created for each host in the cluster.
Windows mount proxy node	<i>hostname_HV_MP_WIN</i>	One of two nodes in a mount proxy node pair that is required for mount operations for the file restore interface. For clusters, a Windows mount proxy node is created for each host in the cluster.
Linux mount proxy node	<i>hostname_HV_MP_LNX</i>	One of two nodes in a mount proxy node pair that is required for mount operations for the file restore interface. For clusters, a Linux mount proxy node is created for each host in the cluster.

You can also add a prefix and suffix to the default node names, as shown: as *prefix_hostname_HV_TGT_suffix*. For instructions, see [“Customizing node names”](#) on page 20.

Policy management at the virtual machine level

Storage requirements for Hyper-V virtual machine backups are determined by IBM Storage Protect server management classes.

You can set different policies for different virtual machines. Although the default management class determines storage characteristics for all Hyper-V backups, you can override the default management class or specify a management class to use for the Hyper-V control files.

You can change the default management class for Hyper-V virtual machine backups with the `vmmc` option. You can change the default management class for Hyper-V control files with the `vmctlmc` option.

Related reference

[“Vmmc” on page 206](#)

Use the `vmmc` option to store virtual machine backups by using a management class other than the default management class.

[“Vmctlmc” on page 190](#)

This option specifies the management class to use when backing up virtual machine control files.

Incremental forever backup strategy

An incremental forever backup strategy minimizes backup windows while providing faster recovery of your data.

IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V provides a backup strategy called incremental forever. This backup solution requires only one initial full backup. Afterward, an ongoing (forever) sequence of incremental backups occurs. The incremental forever backup solution provides these advantages:

- Reduces the amount of data that goes across the network.
- Reduces data growth because all incremental backups contain only the blocks that changed since the previous backup.
- No comparison with the backup target is needed since only changed blocks are identified.
- Minimizes impact to the client system.
- Reduces the length of the backup window.
- No need to schedule an initial full backup as a separate schedule: the first issue of an incremental forever backup automatically defaults to an incremental forever full backup.

In addition, the restore process is optimized, as only the latest versions of blocks that belong to a restored backup are restored. Since the same area on the production disk is recovered only one time, the same block is not written to multiple times. As a result of these advantages, incremental forever is the preferred backup strategy.

Snapshot management with Windows PowerShell

On a Microsoft Hyper-V system, you can use Windows PowerShell "cmdlets" to remove (undo) snapshots that were created by IBM Storage Protect for a Hyper-V virtual machine.

You can use these cmdlets only on the Hyper-V system. You cannot remove snapshots from the Microsoft System Center Virtual Machine Manager.

Hyper-V systems issue cautionary messages to discourage you from editing virtual hard disks that contain snapshots, or virtual hard disks that are associated with a chain of differencing (incremental-forever) snapshots. Instead, use the cmdlets to manage snapshots to minimize the risk of data loss.

For a list of cmdlets that are available for Hyper-V, go to <http://technet.microsoft.com/en-us/library/hh848559.aspx> and read the information for the available cmdlets. Use the **Get-VMSnapshot** cmdlet with the **-SnapshotType Recovery** parameter to retrieve snapshots that are associated with a virtual machine (VM). Use the **Remove-VMSnapshot** cmdlet to remove a snapshot. Removing a snapshot

merges the information that the snapshot wrote to the snapshot differences file (the AVHDX file) back to the VM hard disk (the VHDX file).

If multiple types of snapshots exist for a VM, you can filter the results by snapshot type when you remove a snapshot. For example, to remove only those snapshots that have the snapshot type of "recovery", run the following cmdlet:

```
get-vmsnapshot * | where snapshottype -eq recovery | remove-vmsnapshot
```

Automated client failover for Data Protection for Hyper-V

If you backed up data to the IBM Storage Protect server, Data Protection for Microsoft Hyper-V can automatically fail over to the 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 Hyper-V 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 Microsoft Hyper-V 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.
- If you want to restore data by using the Data Protection for Microsoft Hyper-V or the IBM Storage Protect file restore interface, automatic failover is not available. 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 or instant access operation by using the **restore vm** command with the **vmrestoretype** option. For more information, see [“Vmrestoretype” on page 209](#).

Limitations on Hyper-V backup operations

Before you start a Hyper-V backup operation, review the limitations. Some limitations apply to all Hyper-V backup operations, while others apply only to Hyper-V backups on Windows Server 2012 or 2012 R2 or Windows Server 2016 environments.

Limitations that apply to all Hyper-V backups

You cannot run concurrent backup or restore operations on the same host. For example, if you run two or more **backup vm** or **restore vm** commands on the same host at the same time, one of the backup or restore operations fails with an error message. Starting in Data Protection for Microsoft Hyper-V 8.1.6, the Data Protection for Microsoft Hyper-V Management Console will queue up backup and restore tasks that are submitted to the same host. Only one backup or restore task will be active on a host, and additional backup or restore tasks will be in the pending state until the active task completes. At that time, the next pending task will become active.

Data Protection for Microsoft Hyper-V supports incremental-forever full backup and incremental-forever incremental backup operations for Microsoft Hyper-V virtual machines (VMs) in VHDX disk format only. If you need to back up Hyper-V VMs in VHD disk format, use the version 7.1 backup-archive client (without Data Protection for Microsoft Hyper-V) to create an image backup of the full VM. Issue the version 7.1 backup-archive client command **dsmc backup vm vmname -vmbackuptype=hypervfull -mode=full** to create an image backup of all objects on a Microsoft Hyper-V virtual machine VHD or VHDX disk. Optionally, convert .vhd files to .vhdx format according to instructions available in Microsoft documentation.

Data Protection for Microsoft Hyper-V support for VM backup operations is limited to VM names and Hyper-V host or cluster names that contain English 7-bit ASCII characters only. VM names and Hyper-V host or cluster names that use other language characters are not currently supported. More character restrictions are listed in [“Unsupported characters in virtual machine and Hyper-V host or cluster names”](#) on page 231.

The Microsoft Windows Management Instrumentation (WMI) Service (**winmgmt**) must be running on the systems where Data Protection for Microsoft Hyper-V, IBM Storage Protect backup-archive client, and IBM Storage Protect recovery agent are installed. Operations fail if the WMI Service is not running. Therefore, do not turn off the WMI Service.

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.

You cannot back up a VM with a shared virtual hard disk.

Snapshot differential backup operations are not supported in the Hyper-V 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 Microsoft Hyper-V data mover is also installed.

Limitations that apply only to VM backups on Windows Server 2012 and 2012 R2

Data Protection for Microsoft Hyper-V does not back up VMs with attached physical disks (pass-through disks such as iSCSI disks). This limitation occurs because Data Protection for Microsoft Hyper-V uses Volume Shadow Copy Service (VSS) for backup operations and VSS cannot create a snapshot of the physical disks. If you try to back up a VM with attached physical disks, the backup operation of the VM with the physical disk fails, but backup operations continue for other VMs.

Hyper-V configurations on the Windows Server 2012 R2 operating system are not compatible with Windows Server 2012. As a result, a restore operation from Windows Server 2012 R2 to Windows Server 2012 fails. However, a restore operation from Windows Server 2012 to Windows Server 2012 R2 succeeds. For more information, go to the Microsoft Knowledge Base and search for Article 2868279.

You cannot run multiple backup sessions per VM disk on the Windows Server 2012 operating system. However, you can run multiple backup sessions per VM disk on all other supported Windows operating systems (Windows Server 2012 R2, Windows Server 2016, and later operating systems). You can also run

multiple restore operations per VM disk on all supported Windows operating systems, including Windows Server 2012.

Limitations that apply only to VM backups on Windows Server 2016 and later

You cannot run a VM backup operation during a rolling upgrade of a Hyper-V cluster operating system.

If Data Protection for Microsoft Hyper-V is unable to retrieve the change tracking information, an incremental-forever full backup is run.

Data Protection for Microsoft Hyper-V cannot create an application-consistent snapshot of a VM that is in the Paused state. Only a crash-consistent snapshot can be created of a VM in the Paused state. For example, set the following option in the `dsm.opt` file:

```
INCLUDE.VMSNAPSHOTATTEMPTS VM_name 1 1
```

For Windows Server 2016 only: You cannot install Data Protection for Microsoft Hyper-V on Nano Server for Windows Server 2016. However, you can use Data Protection for Microsoft Hyper-V on Windows Server 2016 to create crash-consistent backups of Nano Server guest VMs.

For late-breaking updates about known issues and limitations, see [technote 1993768](#).

Documentation resources

IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V software provides several components to assist with protecting your virtual machines. As a result, multiple documentation resources are provided to assist with specific tasks.

Table 2. Data Protection for Microsoft Hyper-V documentation resources		
Documentation	Contents	Location
<i>IBM Storage Protect for Virtual Environments Data Protection for Microsoft Hyper-V Installation and User's Guide</i>	Overview information, strategy planning, installation, configuration, back up and restore scenarios, and command-line reference.	IBM Documentation at https://www.ibm.com/support/knowledgecenter/SSERB6/landing/welcome_sserb6.html

Table 2. Data Protection for Microsoft Hyper-V documentation resources (continued)

Documentation	Contents	Location
Online help for the Data Protection for Microsoft Hyper-V Management Console GUI	Back up and restore tasks related to Hyper-V guest virtual machines, configuration, backup management, and backup monitoring.	<p>Start the Data Protection for Microsoft Hyper-V Management Console using either of the following methods:</p> <ul style="list-style-type: none"> On the Windows system, click Start > IBM Storage Protect > DP for Hyper-V Management Console. Open an Administrator command prompt window and enter the following command: <pre>"C:\Program Files\IBM\StorageProtect\DPHyperV\DpHv.msc"</pre> <p>Access the help using either of the following methods:</p> <ul style="list-style-type: none"> In the main page of the Data Protection for Microsoft Hyper-V Management Console, click the links in the Getting Started section. In the menu bar, click Help > Help on Data Protection for Microsoft Hyper-V.
Online help for the IBM Storage Protect file restore interface	Restore individual files and folders from a VM backup.	<p>Start the file restore interface with the URL that is provided by the file restore administrator.</p> <p>Access the help by clicking Help > Product documentation.</p>

Table 2. Data Protection for Microsoft Hyper-V documentation resources (continued)

Documentation	Contents	Location
Online help for the data mover command-line client	Back up and restore tasks related to Hyper-V guest virtual machines.	<p>Start the data mover command-line client by using either of the following methods:</p> <ul style="list-style-type: none"> • On the Windows system, go to Start > IBM Storage Protect > Backup-Archive Command Line. • Open an Administrator command prompt window and change to the backup-archive client installation directory (cd "C:\Program Files\tivoli\tsm\baclient"). Run dsmc.exe. <p>Access the help by using either of the following methods:</p> <ul style="list-style-type: none"> • After you start the command-line client, at the Protect> prompt, enter help to display the table of contents for the help. • To display the help in its own window, open an Administrator command prompt window and change to the backup-archive client installation directory (cd "C:\Program Files\tivoli\tsm\baclient"). <p>Run dsmc.exe help to display the help table of contents. You can also append a topic title to the command to display help for a topic. For example, dsmc help options displays the help topic that describes how to use client options; dsmc help backup vm displays the help for the backup vm command.</p>

Chapter 2. Installing and upgrading Data Protection for Microsoft Hyper-V

Installation of IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V includes planning, installation, and upgrade tasks.

Planning to install Data Protection for Microsoft Hyper-V

Before you install Data Protection for Microsoft Hyper-V, understand the features that are installed and review the system requirements.

Features that are installed

All features for Data Protection for Microsoft Hyper-V are part of the installation suite.

The following components are installed in a typical installation of Data Protection for Microsoft Hyper-V:

- IBM Storage Protect data mover
- Data Protection for Microsoft Hyper-V Management Console
- IBM Storage Protect file restore feature
- Data Protection for Microsoft Hyper-V PowerShell cmdlets
- IBM Storage Protect recovery agent
- IBM Storage Protect web server
- IBM Storage Protect Java™ Virtual Machine (JVM)

You do not have to install any of these features and their support packages separately. For installation instructions, see [“Running a typical installation of Data Protection for Microsoft Hyper-V” on page 24](#).

If you want to install only the data mover for in-guest application protection restore operations, see [“Installing only the Data Protection for Microsoft Hyper-V data mover” on page 28](#). The recovery agent is included in the data mover installation. You can no longer install the IBM Storage Protect recovery agent separately.

If you want to remotely manage Data Protection for Microsoft Hyper-V, install only the Data Protection for Microsoft Hyper-V Management Console on a separate Windows host. For more information, see [“Installing only the Data Protection for Microsoft Hyper-V Management Console” on page 26](#). The PowerShell cmdlets are part of this installation.

You can also install Data Protection for Microsoft Hyper-V on Hyper-V hosts on Windows Server operating systems that were installed with the Server Core option. You can then install Data Protection for Microsoft Hyper-V Management Console on another Windows Server or Windows 10 client to remotely manage Data Protection for Microsoft Hyper-V. For more information, see [“Installing and configuring Data Protection for Microsoft Hyper-V on Windows Server Core systems” on page 32](#).

Determining system requirements

Data Protection for Microsoft Hyper-V requires a minimum amount of hardware, disk space, memory, and software.

The following table describes the minimum hardware requirements that are needed to install Data Protection for Microsoft Hyper-V.

Table 3. Minimum hardware requirements for Data Protection for Microsoft Hyper-V		
Component	Minimum requirement	Preferred
System	x64 processor	Not applicable

<i>Table 3. Minimum hardware requirements for Data Protection for Microsoft Hyper-V (continued)</i>		
Component	Minimum requirement	Preferred
Memory	4 GB RAM	16 GB RAM
Available hard disk	2 GB	3.5 GB
NIC Card	1 NIC - 100 Mbps	1 NIC - 10 Gbps

Data Protection for Microsoft Hyper-V requires the Hyper-V role to be installed on the Microsoft Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 operating system. The Hyper-V Server, a stand-alone product that contains only the Windows hypervisor, is also supported.

To ensure robustness and performance on a Windows Server 2012 or 2012 R2 system, use a VSS hardware provider rather than a software provider.

For Windows Server 2016 only: You cannot install Data Protection for Microsoft Hyper-V on Nano Server for Windows Server 2016. However, you can use Data Protection for Microsoft Hyper-V on Windows Server 2016 to create crash-consistent backups of Nano Server guest VMs.

For application consistent backups and file restore operations on Linux systems, Linux Integration Services (LIS) must be installed on Linux virtual machines (VMs) and the Linux mount proxy. For more information, see [technote 0880755](#).

For detailed software and hardware requirements for Data Protection for Microsoft Hyper-V, see the Version 8.1.27 section in [technote 1505139](#).

For detailed software requirements for application protection of VMs that host Microsoft Exchange Server or Microsoft SQL Server, see [technote 2017347](#).

For prerequisites for the file restore feature, see [“File restore prerequisites” on page 82](#).

Required communication ports

Before you install Data Protection for Microsoft Hyper-V, ensure that specific communication ports are open in the firewall.

The following TCP ports are used by Data Protection for Microsoft Hyper-V. These ports must be open on each computer's respective firewall.

<i>Table 4. Required communication ports for Data Protection for Microsoft Hyper-V</i>			
Computer	Function	Inbound TCP Ports	Outbound TCP Ports
Hyper-V host	All	1581, 1582, 3260, 9081	135, 445, 1500, 1581, 9081
Windows VM	File restore, application protection	135, 445	Does not apply
Linux mount proxy	File restore	1581	22, 1581, 3260
Linux VM	File restore	22	Does not apply

The following table shows what ports are used by what components.

<i>Table 5. Communication ports that are used by components</i>	
Component	TCP Ports
SSH	22
WMI	135, 445

Table 5. Communication ports that are used by components (continued)	
Component	TCP Ports
IBM Storage Protect server	1500
Client acceptor (CAD)	1581, 1582
iSCSI	3260
REST API	9081

Restriction: The Windows mount proxy on the Hyper-V host and the Linux mount proxy must be on the same subnet to support iSCSI traffic.

If any of these ports are changed during the configuration, the firewall rules must be updated.

Upgrading Data Protection for Microsoft Hyper-V

Review the tasks that you need to do when you upgrade to Data Protection for Microsoft Hyper-V Version 8.1.27 from a previous version.

Compatibility with different versions

In environments with multiple clusters and hosts, Data Protection for Microsoft Hyper-V is compatible with later versions.

When you deploy Data Protection for Microsoft Hyper-V to multiple clusters and hosts within your environment, the installed product versions are compatible with later versions. Specifically, when new versions of Data Protection for Microsoft Hyper-V are introduced into your environment, earlier versions of the Data Protection for Microsoft Hyper-V Management Console and PowerShell cmdlets can connect to the newer versions. This compatibility gives you time to update all Data Protection for Microsoft Hyper-V deployments in the environment to the latest levels.

However, a Data Protection for Microsoft Hyper-V Management Console or PowerShell cmdlet cannot connect to an earlier version of Data Protection for Microsoft Hyper-V. A message prompts you to either upgrade the earlier deployment to a newer version, or to use the management console or PowerShell cmdlet that is provided with the earlier deployment.

Example

The following table illustrates the compatibility between previous versions of Data Protection for Microsoft Hyper-V (version 8.1.4 and version 8.1.6) and Data Protection for Microsoft Hyper-V version 8.1.27, in environments where Data Protection for Microsoft Hyper-V is deployed to multiple clusters or hosts.

In this example, Data Protection for Microsoft Hyper-V version 8.1.4 or version 8.1.6 and Data Protection for Microsoft Hyper-V version 8.1.27 are installed in different clusters in the environment.

Table 6. Compatibility examples		
Data Protection for Microsoft Hyper-V Management Console or cmdlet version	Data Protection for Microsoft Hyper-V version	Compatible?
version 8.1.4 or version 8.1.6	version 8.1.27	Yes. All operations work as if you are connected to a version 8.1.4 or version 8.1.6 deployment.

Table 6. Compatibility examples (continued)		
Data Protection for Microsoft Hyper-V Management Console or cmdlet version	Data Protection for Microsoft Hyper-V version	Compatible?
version 8.1.27	version 8.1.4 or version 8.1.6	No. Upgrade the version 8.1.4 or version 8.1.6 deployment to the newest version, or use the management console or Power Shell cmdlets that are provided with the earlier deployment.

Renaming nodes on the IBM Storage Protect server

Before you upgrade your environment from Data Protection for Microsoft Hyper-V version 8.1.2 or earlier to version 8.1.27, you and the IBM Storage Protect server administrator must rename the nodes on the server.

About this task

When you rename the exiting node names on the IBM Storage Protect, you must use the naming convention as described in Step “1” on page 18 .

Restriction: If you use the configuration wizard to configure Data Protection for Microsoft Hyper-V, you must complete the configuration before you can restore older virtual machine (VM) backups that were created with Data Protection for Microsoft Hyper-V version 8.1.2 or earlier. Otherwise, you cannot restore older VM backups with the Data Protection for Microsoft Hyper-V Management Console.

If you manually configure Data Protection for Microsoft Hyper-V and use the data mover command line to restore VMs, the older node names are still operational until you run the configuration wizard.

Procedure

The IBM Storage Protect server administrator completes the following steps:

1. Use the **RENAME NODE** server command to rename the existing Hyper-V node name (specified by the `asnodename` option) to a new target node name that conforms to the following naming conventions:
 - For a stand-alone Hyper-V host environment: `hostname_HV_TGT`
 - For a cluster environment: `clustername_HV_TGT`

For example, for a cluster with cluster node name `Cluster1`, the new target node name becomes `Cluster1_HV_TGT` or `prefix_Cluster1_HV_TGT_suffix`.

You can also add a prefix and a suffix to the default node name. For example, `prefix_hostname_HV_TGT_suffix` or `prefix_clustername_HV_TGT_suffix`.

For instructions about adding a prefix and suffix to the node name, see “Customizing node names” on page 20.

Restriction: You cannot use node names that do not conform to these naming conventions. When you run the Data Protection for Microsoft Hyper-V configuration wizard, the new target node and associated data mover nodes with the new naming conventions are automatically registered on the IBM Storage Protect server. The necessary Windows services are also configured on the local Windows host.

2. Use the **UPDATE SCHEDULE** server command to update existing schedules with the following required parameters:
 - Include the `ACTION=BACKUP` and `SUBACTION=VM` parameters in the schedule definition.
 - Update the option string as follows:

- For a stand-alone host name: `options=' - asnodename=hostname_HV_TGT -domain.vmfull="all-vm" ' or options=' - asnodename=prefix_hostname_HV_TGT_suffix -domain.vmfull="all-vm" '`
- For a cluster name: `options=' - asnodename=clustername_HV_TGT -domain.vmfull="all-vm" ' or options=' - asnodename=prefix_clustername_HV_TGT_suffix -domain.vmfull="all-vm" '`

For more information, see [“Managing backup schedules for a host or cluster machine”](#) on page 72.

- Optional: Update the node replication parameters by issuing the `REPLICATE NODE` command on the IBM Storage Protect server:
 - For a stand-alone host, replicate data on the `hostname_HV_TGT` or `prefix_hostname_HV_TGT_suffix` node.
 - For a cluster, replicate data on the `clustername_HV_TGT` or `prefix_clustername_HV_TGT_suffix` node.

Complete the following tasks on a Hyper-V host:

- Upgrade Data Protection for Microsoft Hyper-V to version 8.1.27 on a stand-alone host or on all hosts in a cluster.

For instructions, see [“Running a typical installation of Data Protection for Microsoft Hyper-V”](#) on page 24.

- Run the configuration wizard on the Hyper-V host. For clusters, run the wizard on one of the hosts in the cluster, typically your local Windows host.

For instructions, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard”](#) on page 39.

- Associate a schedule with the new target node name by using the **Backup Management** window in the Data Protection for Microsoft Hyper-V Management Console.

For instructions, see [“Managing backup schedules for a host or cluster machine”](#) on page 72.

- Verify your configuration by running backup and restore operations in the Data Protection for Microsoft Hyper-V Management Console.

For instructions, see:

- [“Running an ad hoc backup of a virtual machine”](#) on page 77
- [“Restoring a virtual machine”](#) on page 78

- Complete the following cleanup tasks after the configuration is verified:

- The IBM Storage Protect server administrator deletes the old data mover nodes by using the **REMOVE NODE** server command.
- The Hyper-V administrator removes the services that were created for the old cluster node and data mover nodes by running the `dsmcutil remove` command on the stand-alone host or each host in a cluster.

For more information, see the **REMOVe** command in [Dsmcutil commands: Required options and examples](#).

Results

You can protect your Hyper-V VMs with Data Protection for Microsoft Hyper-V.

Tips for viewing backup history: The backup history that occurred before the node update is no longer available. However, all backups can still be restored with the **Restore** wizard or the command line. After the node update, the backup history is available for the initial and subsequent backup operations.

Immediately after the node update in cluster configurations, you can view and restore VM backups only from the cluster view, not from the host view. The host view contains only the VMs that are owned by that host node. After the node update, the backups are not owned by the host node. After successful backups are run, the VMs can again be backed up and restored from the host view.

What to do next

In some situations, one or both of the following tasks might need to be completed:

- The IBM Storage Protect server administrator verifies that the target node is granted proxy authority for the data mover node by issuing the **QUERY PROXY NODE** server command.
- The Hyper-V administrator restarts the client acceptor service on the Hyper-V host.

The IBM Storage Protect server administrator runs the schedule so that backup reporting can be displayed correctly for the updated nodes.

Related concepts

“How IBM Storage Protect nodes are used in Data Protection for Microsoft Hyper-V” on page 6
Data Protection for Microsoft Hyper-V communicates to VMs during backup, restore, and mount operations through IBM Storage Protect nodes.

Customizing node names

You can add a prefix, suffix, or both to default node names. In this way, you can customize the node names that are generated automatically by the configuration wizard.

About this task

When you use the configuration wizard to configure Data Protection for Microsoft Hyper-V, the nodes that are created conform to the following default naming conventions:

hostname_HV_TGT (or *clustername_HV_TGT* for clusters)
hostname_HV_DM
hostname_HV_MP_WIN (if the file restore feature is enabled)
hostname_HP_MP_LNX (if the file restore feature is enabled)

However, you can customize the node names. For example, you might have to customize node names to support a multitenant environment in which multiple tenants host their virtual machines on the same server. To differentiate the nodes based on tenant, you can add a prefix, suffix, or both to the default node names.

You can customize the node names for a new Data Protection for Microsoft Hyper-V configuration or for an existing configuration.

Procedure

To customize node names, complete the following steps:

1. Create a text file named `hvConfig.props` in the `C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI` directory on the Hyper-V host in a stand-alone or cluster environment.
2. Edit the `hvConfig.props` file and add the following two statements to the file:

```
node_prefix=prefix
node_suffix=suffix
```

where *prefix* is the text string that you want to add to the beginning of the node name, and *suffix* is the text string that you want to append to the node name.

You can specify only the prefix, only the suffix, or both a prefix and a suffix. The total length of the node name (including the prefix, suffix, or both) cannot exceed 64 characters.

If you leave the text string blank or remove the statement, the default node name remains unchanged. If you do not want to use any prefix or suffix, do not create the `hvConfig.props` file.

The resulting customized node names follow this pattern:

prefix_hostname_HV_TGT_suffix (or *prefix_clustername_HV_TGT_suffix* for clusters)

prefix_hostname_HV_DM_suffix
prefix_hostname_HV_MP_WIN_suffix (if the file restore feature is enabled)
prefix_hostname_HP_MP_LNX_suffix (if the file restore feature is enabled)

3. For a cluster environment, create the
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\hvConfig.props
file on each host in the cluster and specify the same properties on each host.

Important: All hosts in the cluster must have this file before you can run the configuration wizard.

4. If you previously configured Data Protection for Microsoft Hyper-V with the default node names, you must rename the nodes on the IBM Storage Protect server.

For instructions, see [“Renaming nodes on the IBM Storage Protect server” on page 18.](#)

5. Run the configuration wizard on the Hyper-V host. The prefix, suffix, or both are added to the naming convention of the nodes.

Results

For example, you want to add the prefix "SP" and suffix "DEPT1" to the Data Protection for Microsoft Hyper-V node names. You added the following statements to the C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\hvConfig.props file:

```
node_prefix=SP
node_suffix=DEPT1
```

For a stand-alone host named MYHOST, the following nodes are created when you run the configuration wizard (if the file restore feature is enabled):

```
SP_MYHOST_HV_TGT_DEPT1
SP_MYHOST_HV_DM_DEPT1
SP_MYHOST_HV_MP_WIN_DEPT1
SP_MYHOST_HV_MP_LNX_DEPT1
```

If you have a cluster with cluster name MYCLUSTER, hosts HOSTA and HOSTB, the following nodes are created (if the file restore feature is enabled):

```
SP_MYCLUSTER_HV_TGT_DEPT1
SP_HOSTA_HV_DM_DEPT1
SP_HOSTA_HV_MP_WIN_DEPT1
SP_HOSTA_HV_MP_LNX_DEPT1
SP_HOSTB_HV_DM_DEPT1
SP_HOSTB_HV_MP_WIN_DEPT1
SP_HOSTB_HV_MP_LNX_DEPT1
```

What to do next

You can verify the values for the **node_prefix** and **node_suffix** parameters by running the **Get-DpHvHostConfiguration** Windows PowerShell cmdlet.

For example, from a PowerShell prompt, issue the following cmdlet:

```
PS C:\Users\administrator> Get-DpHvHostConfiguration -Session $session
```

Related concepts

[“How IBM Storage Protect nodes are used in Data Protection for Microsoft Hyper-V” on page 6](#)

Data Protection for Microsoft Hyper-V communicates to VMs during backup, restore, and mount operations through IBM Storage Protect nodes.

Upgrade considerations for RCT backups

Before you upgrade to Data Protection for Microsoft Hyper-V Version 8.1.2 or later, review the considerations that apply to virtual machine (VM) backup operations on Windows Server 2016.

- When you upgrade your Hyper-V environment from Windows Server 2012 or 2012 R2 to Windows Server 2016, the VM version of the virtual machines is not updated automatically. The Hyper-V administrator must update the VMs to the new version after the environment is upgraded to Windows Server 2016. Data Protection for Microsoft Hyper-V version 8.1.2 or later does not back up VMs that are not updated to the new VM version.

Ensure that the guest VM is offline before you update the VM version. You can update the VM version in the Hyper-V Manager or the [Update-VMVersion](#) cmdlet.

- VM backup operations with resilient change tracking (RCT) require the Hyper-V VM to be Version 6.2 or later.

Data Protection for Microsoft Hyper-V version 8.1.0 and earlier continues to support earlier VM versions by using the VSS backup method.

Related tasks

[“Migrating from VSS backups to RCT backups” on page 22](#)

To take advantage of the resilient change tracking (RCT) backup feature in Data Protection for Microsoft Hyper-V version 8.1.2 or later, migrate your virtual machine (VM) backup operations from the Microsoft Volume Shadow Copy Service (VSS) to RCT.

Migrating from VSS backups to RCT backups

To take advantage of the resilient change tracking (RCT) backup feature in Data Protection for Microsoft Hyper-V version 8.1.2 or later, migrate your virtual machine (VM) backup operations from the Microsoft Volume Shadow Copy Service (VSS) to RCT.

Before you begin

- Verify that the Hyper-V VM is at version 6.2 or later. You can determine the VM version in the Hyper-V Manager or by running the [Get-VM](#) cmdlet.
- When you migrate your Hyper-V environment from Windows Server 2012 or 2012 R2 to Windows Server 2016, the VM version of the Hyper-V VMs is not updated automatically. You must update the VMs to the new version before they can be backed up by Data Protection for Microsoft Hyper-V.

Ensure that you take the guest VM offline before you update the VM version of a VM. You can update the VM version in the Hyper-V Manager or with the [Update-VMVersion](#) cmdlet.

Procedure

To migrate VSS backups to RCT:

1. Install and configure Data Protection for Microsoft Hyper-V version 8.1.27 on the Hyper-V host server on the Windows Server 2016 operating system.
2. Run an incremental-forever full backup operation on your VMs.

All Data Protection for Microsoft Hyper-V backup operations in the Windows Server 2016 or later environment use RCT backups.

Results

- Because previous VSS backups do not have RCT change-tracking information, an incremental-forever full backup is created the first time you back up a VM with Data Protection for Microsoft Hyper-V version 8.1.27.

- VSS backups are disabled after the initial backup of a VM with RCT.
- With Data Protection for Microsoft Hyper-V version 8.1.27, you can still restore VMs that were backed up on Windows Server 2016 in version 8.1.0. Subsequent backups of VMs use RCT.

Related concepts

[“Virtual machine backups with resilient change tracking \(RCT\)” on page 2](#)

For Hyper-V backups on Microsoft Windows Server 2016 or later versions, the resilient change tracking (RCT) feature is used to back up virtual machines (VMs).

[“Virtual machine backups with Volume Shadow Copy Service \(VSS\)” on page 1](#)

For Hyper-V backups on Windows Server 2012 and 2012 R2, Microsoft Volume Shadow Copy Service (VSS) is used to create consistent snapshots of virtual machines (VMs) during backup operations.

Upgrading Data Protection for Microsoft Hyper-V by using the installation wizard

You can upgrade to Data Protection for Microsoft Hyper-V Version 8.1.27 by using the installation wizard.

Before you upgrade Data Protection for Microsoft Hyper-V, review the following information:

- Before the upgrade, **do not** uninstall the previous version. If you upgrade Data Protection for Microsoft Hyper-V by first uninstalling the previous version and then installing the new version, the upgrade is treated as a new installation. As a result, when you initially start the Data Protection for Microsoft Hyper-V Management Console after the installation, the configuration wizard will be opened automatically and you will have to go through the configuration again.
- Data Protection for Microsoft Hyper-V can be installed only in the default location (C:\Program Files\IBM\StorageProtect). If you previously installed Data Protection for Microsoft Hyper-V in a custom location, the program files from the previous version, except for the log and trace files, are moved to the default installation location during the upgrade.
- If you are upgrading from a previous version of Data Protection for Microsoft Hyper-V that was installed in a custom location, ensure that you stop all IBM Storage Protect client acceptor and agent services before you upgrade to Data Protection for Microsoft Hyper-V version 8.1.27. The services are typically called **TSM Client Acceptor** and **TSM Remote Client Agent** in the **Services** window. Stopping these services before the upgrade will ensure that there is no extended delay during the upgrade process.

To upgrade to Data Protection for Microsoft Hyper-V version 8.1.27, follow the instructions in [“Installing Data Protection for Microsoft Hyper-V components” on page 23](#), but with the following restriction: In the installation wizard, you must select the same installation type that was used to install the previous version of Data Protection for Microsoft Hyper-V. For example, if you selected **Typical Installation** to install the previous version, you must also select **Typical Installation** during the upgrade.

Installing Data Protection for Microsoft Hyper-V components

Run a typical installation to install all of the Data Protection for Microsoft Hyper-V components. You can then install separate components as needed for your use case.

Tip: You can use the same instructions to upgrade Data Protection for Microsoft Hyper-V. However, review the information in [“Upgrading Data Protection for Microsoft Hyper-V” on page 17](#) before you begin the upgrade.

Download and extract the installation package

Before you install Data Protection for Microsoft Hyper-V, you must download the installation package and extract the installation files from the package.

Before you begin

For the most recent information, updates, and maintenance fixes, go to [IBM Storage Protect for Virtual Environments - IBM Support](#).

Procedure

1. Download the Data Protection for Microsoft Hyper-V package from [IBM Passport Advantage®](#) or [Fix Central](#).
2. Extract the compressed installation file that you downloaded:
 - a) Copy the downloaded compressed installation package to a local disk or to a network-accessible share. Be sure to extract the installation files to an empty directory (*extract_folder*).
 - b) To extract the installation files to the same directory, double-click the compressed installation package.

By default, the uncompressed files are stored in the current disk drive, in the *extract_folder*\TSMHYPERV_WIN directory.

If the installation program detects files from another Data Protection for Microsoft Hyper-V installation attempt in this directory, you are prompted to specify whether to overwrite the old files. If you see a prompt about overwriting files, enter (A) lways to overwrite the existing files; this selection ensures that only the files from the current installation are used.

Results

The Data Protection for Microsoft Hyper-V installation program (spinstall.exe) is located in the *extract_folder*\TSMHYPERV_WIN directory.

What to do next

[Install Data Protection for Microsoft Hyper-V.](#)

Installing Data Protection for Microsoft Hyper-V by using the installation wizard

Use the installation wizard to complete a typical installation of Data Protection for Microsoft Hyper-V or to install the available components separately.

Running a typical installation of Data Protection for Microsoft Hyper-V

Instructions are provided for a typical installation of the IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V software.

Before you begin

If you are in a cluster environment, ensure that you install the Data Protection for Microsoft Hyper-V package on each host in the cluster.

On each host where Data Protection for Microsoft Hyper-V is installed, ensure that the HTTPS port that is used to communicate with Data Protection for Microsoft Hyper-V is open in the firewall. Unless specified otherwise, the default port number of 9081 is used.

If you are upgrading from Data Protection for Microsoft Hyper-V Version 8.1.2 or earlier, complete the tasks in [“Renaming nodes on the IBM Storage Protect server”](#) on page 18.

If you are upgrading Data Protection for Microsoft Hyper-V, review the information in [“Upgrading Data Protection for Microsoft Hyper-V”](#) on page 17.

Ensure that you downloaded and extracted the installation package as described in [“Download and extract the installation package”](#) on page 23.

About this task

A typical installation includes all the features of Data Protection for Microsoft Hyper-V, including the data mover, the Data Protection for Microsoft Hyper-V Management Console, the PowerShell cmdlets, and the IBM Storage Protect recovery agent.

To install only the Data Protection for Microsoft Hyper-V Management Console for remote management, see [“Installing only the Data Protection for Microsoft Hyper-V Management Console” on page 26](#). The PowerShell cmdlets are part of this installation.

To install only the data mover for in-guest application protection restore operations, see [“Installing only the Data Protection for Microsoft Hyper-V data mover” on page 28](#). The recovery agent is included in the data mover installation.

Restriction: The Data Protection for Microsoft Hyper-V installer automatically disables the automount feature with the **diskpart** command on the Windows operating system. This action is required to show correct drive letter assignments and to hide the system reserved disk in the IBM Storage Protect file restore interface.

If you do not plan to run file restore operations, or if you do not care that incorrect drive letter assignments and the system reserved disk are displayed in the file restore interface, you can enable the automount feature after the installation is completed.

Procedure

Complete the following steps on a single Hyper-V host or on each host in a cluster:

1. [“Download and extract the installation package” on page 23](#).
2. To start the installation program, double-click the `spinstall.exe` file. Choose the language for the installation process, and then click **Next**.
3. On the **Welcome to the InstallShield Wizard for IBM Storage Protect for Virtual Environments: Data Protection for Hyper-V suite** page, click **Next**.
4. On the **License Agreement** page, read the terms of the license agreement. Click **I accept the terms in the license agreement** and click **Next**. If you do not accept the terms of the license agreement, the installation ends and you must click **Cancel** to exit the installation wizard.
5. On the **Installation Destination Folder** page, note the installation location and click **Next**.

Data Protection for Microsoft Hyper-V can be installed only in the default location (C:\Program Files\IBM\StorageProtect). If you previously installed Data Protection for Microsoft Hyper-V in a custom location, the program files from the previous version, except for the log and trace files, are moved to the default installation location during the upgrade.

6. On the **Installation Type** page, click **Typical Installation**. The installation process begins immediately. You cannot change your selection after the installation process begins.

Tip: The installation process might take several minutes to complete while the Data Protection for Microsoft Hyper-V, JVM, data mover, web server, framework, and recovery agent packages are being installed.

7. On the **Install Wizard Completed** page, click **Finish** to exit the installation wizard.

The Data Protection for Microsoft Hyper-V Management Console starts immediately after the wizard is closed.

If you do not want to start the configuration wizard now, clear the **Launch Data Protection for Microsoft Hyper-V Management Console now** box and click **Finish** to exit the wizard.

Results

Data Protection for Microsoft Hyper-V is installed.

The following installed components appear in the **Programs and Features** control panel in the Windows operating system:

- IBM Storage Protect Client

- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V License
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite
- IBM Storage Protect for Virtual Environments: Framework
- IBM Storage Protect for Virtual Environments: Recovery agent
- IBM Storage Protect JVM
- IBM Storage Protect WebServer

What to do next

Before you attempt a backup or restore operation, or use the file restore interface, complete the tasks in [Configure Data Protection for Microsoft Hyper-V with the wizard](#).

Before you attempt to mount the backup of a Hyper-V virtual machine disk to restore a file, complete the tasks in [“Configuring the IBM Storage Protect recovery agent GUI ” on page 50](#).

Installing only the Data Protection for Microsoft Hyper-V Management Console

You can install only the Data Protection for Microsoft Hyper-V Management Console on a Windows host for the remote management of Data Protection for Microsoft Hyper-V.

Before you begin

Ensure that the HTTPS port that is used to communicate with Data Protection for Microsoft Hyper-V Management Console is open in the firewall. The default port is 9081, unless you are using a different port. For more information, see [“Required communication ports” on page 16](#).

Ensure that you downloaded and extracted the installation package as described in [“Download and extract the installation package” on page 23](#).

About this task

This installation includes only the Data Protection for Microsoft Hyper-V Management Console, Data Protection for Microsoft Hyper-V PowerShell cmdlets, and the Data Protection for Microsoft Hyper-V license file.

Restriction: The Data Protection for Microsoft Hyper-V installer automatically disables the automount feature with the **diskpart** command on the Windows operating system. This action is required to show correct drive letter assignments and to hide the system reserved disk in the IBM Storage Protect file restore interface.

If you do not plan to run file restore operations, or if you do not care that incorrect drive letter assignments and the system reserved disk are displayed in the file restore interface, you can enable the automount feature after completing the installation.

Procedure

Complete the following steps on a Windows computer that you want to use to remotely manage Data Protection for Microsoft Hyper-V.

1. [“Download and extract the installation package” on page 23](#).
2. To start the installation program, double-click the `spinstall.exe` file. Choose the language for the installation process, and then click **Next**.
3. On the **Welcome to the InstallShield Wizard for IBM Storage Protect for Virtual Environments: Data Protection for Hyper-V suite** page, click **Next**.

4. On the **License Agreement** page, read the terms of the license agreement. Click **I accept the terms in the license agreement** and click **Next**. If you do not accept the terms of the license agreement, the installation ends and you must click **Cancel** to exit the installation wizard.
5. On the **Installation Destination Folder** page, note the installation location and click **Next**.

Data Protection for Microsoft Hyper-V can be installed only in the default location (C:\Program Files\IBM\StorageProtect). If you previously installed Data Protection for Microsoft Hyper-V in a custom location, the program files from the previous version, except for the log and trace files, are moved to the default installation location during the upgrade.

6. On the **Installation Type** page, click **Advanced Installation**.
7. On the **Advanced Installation** page, click **Install the Data Protection for Microsoft Hyper-V Management Console only**. The installation process begins immediately. You cannot change your selection after the installation process begins.

Tip: The installation process might take several minutes to complete while the necessary packages are being installed.

8. On the **Install Wizard Completed** page, click **Finish** to exit the installation wizard.

The Data Protection for Microsoft Hyper-V Management Console starts immediately after the wizard is closed.

If you do not want to start the configuration wizard now, clear the **Launch Data Protection for Microsoft Hyper-V Management Console now** box and click **Finish** to exit the wizard.

Results

Data Protection for Microsoft Hyper-V Management Console is installed.

The following installed components appear in the **Programs and Features** control panel in the Windows operating system:

- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V License
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite

What to do next

Configure the Data Protection for Microsoft Hyper-V Management Console by completing the following steps:

1. If the configuration wizard does not open automatically, start the Data Protection for Microsoft Hyper-V Management Console by clicking **Start > IBM Storage Protect > DP for Hyper-V Management Console**.
2. In the **Connect to Data Protection for Hyper-V** window, enter the host name and credentials of the stand-alone host or host in the cluster that you want to manage.
3. Complete the tasks in [Configure Data Protection for Microsoft Hyper-V with the wizard](#).

You can also specify the preferred host to log on to by using the **Set-DpHvMmcLoginPreferences** cmdlet. For more information, see [Chapter 8, “Protecting virtual machines by using Windows PowerShell cmdlets,”](#) on page 133.

Related tasks

[“Installing and configuring Data Protection for Microsoft Hyper-V on Windows Server Core systems”](#) on page 32

You can install and configure Data Protection for Microsoft Hyper-V on Hyper-V hosts on Windows Server operating systems that were installed with the Server Core option.

Installing only the Data Protection for Microsoft Hyper-V data mover

You can install the Data Protection for Microsoft Hyper-V data mover for virtual machine (VM) backup and restore operations and in-guest application protection restore operations. This installation also installs the Windows mount proxy for file restore operations.

Before you begin

- Ensure that communication ports are open in the firewall. For the list of ports that need to be open, see [“Required communication ports”](#) on page 16.
- Ensure that you downloaded and extracted the installation package as described in [“Download and extract the installation package”](#) on page 23.
- If you are installing the data mover to protect in-guest applications, ensure that you follow the instructions in the following topics before you install the data mover:
 - [“Installing and configuring software for application protection of Microsoft Exchange Server”](#) on page 87
 - [“Installing and configuring software for application protection of Microsoft SQL Server”](#) on page 103

About this task

The data mover installation includes the data mover, which is used for VM backup and restore operations and in-guest application protection restore operations. This installation also includes the mount proxy for file restore operations. The recovery agent is also included in the installation.

Restriction: The Data Protection for Microsoft Hyper-V installer automatically disables the automount feature with the **diskpart** command on the Windows operating system. This action is required to show correct drive letter assignments and to hide the system reserved disk in the IBM Storage Protect file restore interface.

Procedure

Complete the following steps on the Windows mount proxy machine or the guest VM that is hosting application data:

1. [“Download and extract the installation package”](#) on page 23.
2. To start the installation program, double-click the `spinstall.exe` file. Choose the language for the installation process, and then click **Next**.
3. On the **Welcome to the InstallShield Wizard for IBM Storage Protect for Virtual Environments: Data Protection for Hyper-V suite** page, click **Next**.
4. On the **License Agreement** page, read the terms of the license agreement. Click **I accept the terms in the license agreement** and click **Next**. If you do not accept the terms of the license agreement, the installation ends and you must click **Cancel** to exit the installation wizard.
5. On the **Installation Destination Folder** page, note the installation location and click **Next**.

Data Protection for Microsoft Hyper-V can be installed only in the default location (C:\Program Files\IBM\StorageProtect). If you previously installed Data Protection for Microsoft Hyper-V in a custom location, the program files from the previous version, except for the log and trace files, are moved to the default installation location during the upgrade.

6. On the **Installation Type** page, click **Advanced Installation**.
7. On the **Advanced Installation** page, click **Install the data mover feature or mount proxy**. The installation process begins immediately. You cannot change your selection after the installation process begins.

Tip: The installation process might take several minutes to complete while the necessary packages are being installed.

8. On the **Install Wizard Completed** page, click **Finish** to exit the installation wizard.

Results

The Data Protection for Microsoft Hyper-V data mover is installed.

The following installed components appear in the **Programs and Features** control panel in the Windows operating system:

- IBM Storage Protect Client
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V License
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite
- IBM Storage Protect for Virtual Environments: Framework
- IBM Storage Protect for Virtual Environments: Recovery agent
- IBM Storage Protect JVM
- IBM Storage Protect WebServer

What to do next

For more information about installing and configuring the software for application protection, see one of the following topics:

- [“Installing and configuring software for application protection of Microsoft Exchange Server” on page 87](#)
- [“Installing and configuring software for application protection of Microsoft SQL Server” on page 103](#)

Uninstalling Data Protection for Microsoft Hyper-V

The process for uninstalling IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V is the same for a new installation and for an upgraded version.

Before you begin

Restriction: You can uninstall the IBM Storage Protect recovery agent as part of the uninstallation for the IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite, or you can uninstall the recovery agent separately. You must unmount all virtual volumes before uninstalling the IBM Storage Protect recovery agent. Otherwise, these mounted virtual volumes cannot be unmounted when you reinstall the recovery agent the next time.

Procedure

1. Open the **Control Panel** and click **Uninstall a program**.
2. Uninstall the IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite:
 - a) On the **Uninstall or change a program** page, select **IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V suite** and click **Uninstall**.
 - b) In the **Remove the Program** page of the **InstallShield Wizard**, click **Remove**.

Tip: The uninstallation process might take several minutes to complete.

- c) Click **Finish** in the **InstallShield Wizard Completed** page when the uninstallation is completed. Click the **Refresh** icon to refresh the list of programs.
3. Uninstall the Data Protection for Microsoft Hyper-V license:

- a) On the **Uninstall or change a program** page, select **IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V License** and click **Uninstall**.
 - b) Click **Yes** when prompted.
4. Uninstall the IBM Storage Protect web server:
- a) On the **Uninstall or change a program** page, select **IBM Storage Protect Web Server** and click **Uninstall**.
 - b) Click **Yes** when prompted.
5. Uninstall the IBM Storage Protect Java Virtual Machine (JVM):
- a) On the **Uninstall or change a program** page, select **IBM Storage Protect JVM** and click **Uninstall**.
 - b) Click **Yes** when prompted.

What to do next

You must remove the file restore feature separately. For more information, see [“Removing the file restore feature”](#) on page 38.

Installing Data Protection for Microsoft Hyper-V in silent mode

You can install all Data Protection for Microsoft Hyper-V and data mover features silently on a single system. You can also choose to install only the data mover or only the Data Protection for Microsoft Hyper-V Management Console in silent mode.

Before you begin

Ensure that you downloaded and extracted the installation package as described in [“Download and extract the installation package”](#) on page 23.

About this task

Restriction: All features are installed to their default location. You cannot silently install Data Protection for Microsoft Hyper-V and data mover features to a non-default location.

Procedure

1. Open the command prompt as administrator, and enter the following command:

```
cd extract_folder\TSMHYPERV_WIN
```

2. Use one of the following methods to install Data Protection for Microsoft Hyper-V in silent mode:

- To install all Data Protection for Microsoft Hyper-V and data mover features silently, enter the following command:

```
spinstall.exe /silent
```

The following message is displayed the first time that you mount a volume:

```
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?
```

To proceed with the IBM Storage Protect recovery agent operations, enter **Yes** to register the Virtual Volume Driver.

- To install only the data mover in silent mode, enter the following command:

```
spinstall.exe /silent  
ISFeatureInstall=Client,Java,WebServer,Framework,HyperV,RecoveryAgent  
ComponentsToInstallRA="LAP,mount,shell"
```


- To install only the Data Protection for Microsoft Hyper-V Management Console in silent mode, enter the following command:

```
spinstall.exe /silent ISFeatureInstall=HyperV
```

Uninstalling Data Protection for Microsoft Hyper-V in silent mode

You can silently uninstall Data Protection for Microsoft Hyper-V on a Windows operating system.

About this task

You must first uninstall Data Protection for Microsoft Hyper-V with the Suite installer, and then run a Windows PowerShell cmdlet to remove the following components:

- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V License
- IBM Storage Protect WebServer
- IBM Storage Protect Java Virtual Machine (JVM)

Tip: To see a list of installed Data Protection for Microsoft Hyper-V components, issue the following command as administrator from a PowerShell command prompt:

```
Get-WmiObject -class Win32_Product | Where-Object name -match `
  "IBM Storage Protect*" | Format-List Version,InstallDate,Name,InstallLocation
```

The following procedure can also be used to uninstall Data Protection for Microsoft Hyper-V from Windows Server Core systems.

Tip: To issue PowerShell commands, you must run the PowerShell command prompt in Administrative mode.

Procedure

1. If you ran instant access operations, remove any instant access virtual machines that were created.

For instructions, see [“Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console” on page 131](#).

2. Stop any IBM Storage Protect services that are running. The names of the services typically begin with TSM.

For example, you can run the following command from a PowerShell command prompt:

```
stop-service TSM*
```

3. From a command prompt, issue the following command:

```
cd extract_folder\TSMHYPERV_WIN
```

where *extract_folder* specifies the folder where you extracted the Data Protection for Microsoft Hyper-V installation files.

4. To uninstall Data Protection for Microsoft Hyper-V with the Suite installer, issue the following command from a PowerShell command prompt:

```
Start-Process -FilePath ".\spinstall.exe" -ArgumentList "/silent /remove" -wait
```

5. To remove any remaining components, run the following cmdlet from a PowerShell command prompt. Alternatively, include the following cmdlet in a script for automation.

```
$ve = "IBM Storage Protect for Virtual Environments:"
$lic = "$ve Data Protection for Microsoft Hyper-V License"
$ws = "IBM Storage Protect WebServer"
$jvm = "IBM Storage Protect JVM"

Get-WmiObject -class Win32_Product -Filter "Name='$lic'" | Invoke-WmiMethod `
  -Name "Uninstall" | Select-Object ReturnValue
```

```
Get-WmiObject -class Win32_Product -Filter "Name='$ws'" | Invoke-WmiMethod `
-Name "Uninstall" | Select-Object ReturnValue
Get-WmiObject -class Win32_Product -Filter "Name='$jvm'" | Invoke-WmiMethod `
-Name "Uninstall" | Select-Object ReturnValue
```

Results

In case you want to save the configuration information after the uninstall operation is completed, configuration and property files are located in the following folders:

- C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI
- C:\IBM\StorageProtect\webserver\usr\servers\veProfile\frGUI
- C:\Program Files\IBM\StorageProtect\DPHyperV
- C:\Program Files\Tivoli\TSM\baclient

Related tasks

[“Installing Data Protection for Microsoft Hyper-V in silent mode” on page 30](#)

You can install all Data Protection for Microsoft Hyper-V and data mover features silently on a single system. You can also choose to install only the data mover or only the Data Protection for Microsoft Hyper-V Management Console in silent mode.

Installing and configuring Data Protection for Microsoft Hyper-V on Windows Server Core systems

You can install and configure Data Protection for Microsoft Hyper-V on Hyper-V hosts on Windows Server operating systems that were installed with the Server Core option.

Before you begin

On each host where Data Protection for Microsoft Hyper-V is installed, ensure that the HTTPS port that is used to communicate with Data Protection for Microsoft Hyper-V is open in the firewall. Unless specified otherwise, the default port number of 9081 is used.

About this task

Because local user interfaces are not supported on Server Core, you must install Data Protection for Microsoft Hyper-V silently on a stand-alone host or on each host in a cluster.

You must manage Data Protection for Microsoft Hyper-V by using the Data Protection for Microsoft Hyper-V Management Console from another deployment and pointing it to a stand-alone host or a host in a cluster.

Procedure

1. Run a silent installation of Data Protection for Microsoft Hyper-V on a stand-alone host or on all hosts in a cluster.

For instructions, see [“Installing Data Protection for Microsoft Hyper-V in silent mode” on page 30](#).

2. To manage Data Protection for Microsoft Hyper-V remotely, you must separately install the Data Protection for Microsoft Hyper-V Management Console on another Windows Server or Windows 10 operating system.

For instructions, see [“Installing only the Data Protection for Microsoft Hyper-V Management Console” on page 26](#).

3. If the configuration wizard does not open automatically, start the Data Protection for Microsoft Hyper-V Management Console by clicking **Start > IBM Storage Protect > DP for Hyper-V Management Console**.
4. In the **Connect to Data Protection for Hyper-V window**, enter the host name and credentials of the stand-alone host or host in the cluster that you want to manage.

5. [Configure Data Protection for Microsoft Hyper-V with the wizard.](#)

Results

You can use the Data Protection for Microsoft Hyper-V Management Console to remotely manage daily operations of Data Protection for Microsoft Hyper-V in a stand-alone host or cluster on an operating system that was installed with the Server Core option.

What to do next

You can also specify the preferred host to log on to by using the **Set-DpHvMmcLoginPreferences** cmdlet. For more information, see [Chapter 8, “Protecting virtual machines by using Windows PowerShell cmdlets,”](#) on page 133.

Related concepts

[“Managing data with the Data Protection for Microsoft Hyper-V Management Console”](#) on page 61

The Data Protection for Microsoft Hyper-V Management Console provides a single environment to help you manage the daily operations of Data Protection for Microsoft Hyper-V.

Related tasks

[“Configuring non-default port numbers for Data Protection for Microsoft Hyper-V operations”](#) on page 58

If you do not want to use the default port numbers for the Data Protection for Microsoft Hyper-V web server or REST API services, you can configure different port numbers by using Windows PowerShell cmdlets.

Uninstalling Data Protection for Microsoft Hyper-V from Windows Server Core systems

You can uninstall Data Protection for Microsoft Hyper-V from Windows Server Core systems by using silent mode.

For instructions, see [“Uninstalling Data Protection for Microsoft Hyper-V in silent mode”](#) on page 31.

Related tasks

[“Installing and configuring Data Protection for Microsoft Hyper-V on Windows Server Core systems”](#) on page 32

You can install and configure Data Protection for Microsoft Hyper-V on Hyper-V hosts on Windows Server operating systems that were installed with the Server Core option.

Installing the Linux mount proxy feature

Follow the instructions to install the mount proxy feature on Linux guest virtual machines (VMs) for use with file restore operations.

Upgrading the Linux mount proxy feature from an older version

If the mount proxy feature is already installed on the Linux virtual machine, you can upgrade to the Data Protection for Microsoft Hyper-V Version 8.1.27 Linux mount proxy.

Before you begin

Ensure that the version of the Linux mount proxy that you are upgrading to is the same as the version of the Windows mount proxy. For example, if the Windows mount proxy is at version 8.1.27, the Linux mount proxy must also be at version 8.1.27.

If you are upgrading from the version 8.1.4 Linux mount proxy, you must first uninstall the Linux mount proxy. You can uninstall the version 8.1.4 Linux mount proxy by issuing the following commands:

```
rpm -e TIVsm-BACit.x86_64 TIVsm-BA.x86_64
rpm -e TIVsm-APIcit.x86_64 TIVsm-API64.x86_64
rpm -e gskcrypt64.linux.x86_64.rpm gskssl64.linux.x86_64
```

If you are upgrading from version 8.1.6, you do not need to uninstall the Linux mount proxy.

Procedure

- Upgrade the mount proxy feature directly by installing the version 8.1.27 Linux data mover package.

For instructions, see one of the following topics:

- [“Installing the mount proxy feature on Linux systems” on page 34](#)
- [“Installing the Linux mount proxy feature in silent mode” on page 36](#)

What to do next

After the upgrade, you do not need to reset the Linux mount proxy password as long as you do not relaunch the configuration wizard from the Data Protection for Microsoft Hyper-V Management Console on the Hyper-V host, or delete the encrypted password files in the `/etc/adsm` directory.

You also do not need to restart the Linux system after the upgrade. Simply issue the **kill -9** command to stop any existing active **dsmcad** process. Then, restart the **dsmcad** process to start the client acceptor for version 8.1.27.

Installing the mount proxy feature on Linux systems

If you plan to run file restore operations on Linux guest virtual machines (VMs), you must install the mount proxy feature on Linux systems by using the Linux Data Protection for Microsoft Hyper-V data mover package.

Before you begin

If you are upgrading from an older version of the Linux mount proxy, review the information in [“Upgrading the Linux mount proxy feature from an older version” on page 33](#).

About this task

A mount proxy node is required for mount operations for the file restore interface. The mount proxy node enables the file systems on the mounted VM disks of VM backups to be accessible as mount points for file restore operations.

The Linux mount proxy software is part of the Linux data mover package. It is not included in the standard Windows Data Protection for Microsoft Hyper-V installation package. You must download and install the Linux package separately.

Procedure

As root user, install the mount proxy feature by completing the following steps:

1. Download and extract the installation package:
 - a) Download the Linux Data Protection for Microsoft Hyper-V data mover installation package from one of the following websites:

- [Passport Advantage](#)
- [Fix Central](#)

The download package is typically named `8.x.x.x-TSM4HYPERV.tar.gz`. For example, for Version 8.1.27, the package is named `8.1.27.0-TSM4HYPERV.tar.gz`.

Tip: For the latest information, updates, and maintenance fixes, go to the [IBM Support Portal](#).

- b) Copy the Linux data mover package to a location where you want to store the installation files. For example, create the following directory and copy the installation package to the directory:

```
/extract_folder
```

c) Change to the installation package directory. For example:

```
cd /extract_folder
```

d) Extract the installation files from the installation package by issuing the following command:

```
tar -xvzf 8.1.27.0-TSM4HYPERV.tar.gz
```

The installation files are extracted to the CD directory. For example, the installation files are saved to the following directory:

```
/extract_folder/CD
```

2. Change to the directory that contains the installation file. For example, issue the following command:

```
cd /extract_folder/CD/Linux/DataProtectionForHyperV
```

3. Start the Data Protection for Microsoft Hyper-V installation wizard by issuing the following command:

```
./install-Linux.bin
```

4. Select the language for the installation process and click **OK**.

5. Complete each page of the installation wizard as follows.

Page	Action
Welcome	Click Next .
Software License Agreement	Accept software license agreement and click Next .
Located Installation Directory	Review the installation directory (/opt/tivoli/tsm/) and click Next .
Custom	Ensure that Data Protection for Hyper-V Data Mover is checked and click Next .
Preinstallation Summary	Review the installation summary. To proceed with the installation, click Install .
Review this information before you proceed	Click Next .
Installation Complete	Click Done .

Results

Tip: If you do not want to run the installation wizard, you can use the following methods to install the mount proxy feature:

- To install from the console, issue the following command: `./install-Linux.bin -i console`
- To install in silent mode, see [“Installing the Linux mount proxy feature in silent mode” on page 36](#).

What to do next

Configure the Linux mount proxy for file restore operations. For instructions, see [“Configuring the Linux mount proxy for file restore operations” on page 45](#).

Related tasks

[“Uninstalling the mount proxy feature on Linux systems” on page 37](#)

If you no longer need to run file restore operations on Linux guest virtual machines (VMs), you can uninstall the mount proxy feature on the Linux mount proxy system.

Installing the Linux mount proxy feature in silent mode

If you plan to run file restore operations on Linux guest virtual machines (VMs), you must install the mount proxy feature on Linux systems by using the Linux Data Protection for Microsoft Hyper-V data mover package. If you do not want to install the mount proxy feature interactively, you can install it in silent mode.

Before you begin

If you are upgrading from an older version of the Linux mount proxy, review the information in [“Upgrading the Linux mount proxy feature from an older version”](#) on page 33.

About this task

A mount proxy node is required for mount operations for the file restore interface. The mount proxy node enables the file systems on the mounted VM disks of VM backups to be accessible as mount points for file restore operations.

The Linux mount proxy software is part of the Linux data mover package. It is not included in the standard Windows Data Protection for Microsoft Hyper-V installation package. You must download and install the Linux package separately.

Procedure

As root user, completing the following steps on the Linux guest VM:

1. Download and extract the installation package:

- a) Download the Linux Data Protection for Microsoft Hyper-V data mover installation package from one of the following websites:

- [Passport Advantage](#)
- [Fix Central](#)

The download package is typically named `8.x.x.x-TSM4HYPERV.tar.gz`. For example, for Version 8.1.27, the package is named `8.1.27.0-TSM4HYPERV.tar.gz`.

Tip: For the latest information, updates, and maintenance fixes, go to the [IBM Support Portal](#).

- b) Copy the Linux data mover package to a location where you want to store the installation files. For example, create the following directory and copy the installation package to the directory:

```
/extract_folder
```

- c) Change to the installation package directory. For example:

```
cd /extract_folder
```

- d) Extract the installation files from the installation package by issuing the following command:

```
tar -xvzf 8.1.27.0-TSM4HYPERV.tar.gz
```

The installation files are extracted to the CD directory. For example, the installation files are saved to the following directory:

```
/extract_folder/CD
```

2. Change to the directory that contains the installation file. For example, issue the following command:

```
cd /extract_folder/CD/Linux/DataProtectionForHyperV
```

3. Issue the following command to start the installation in silent mode and accept the Data Protection for Microsoft Hyper-V license.

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

What to do next

Configure the Linux mount proxy for file restore operations. For instructions, see [“Configuring the Linux mount proxy for file restore operations”](#) on page 45.

Related tasks

[“Installing the mount proxy feature on Linux systems”](#) on page 34

If you plan to run file restore operations on Linux guest virtual machines (VMs), you must install the mount proxy feature on Linux systems by using the Linux Data Protection for Microsoft Hyper-V data mover package.

Uninstalling the mount proxy feature on Linux systems

If you no longer need to run file restore operations on Linux guest virtual machines (VMs), you can uninstall the mount proxy feature on the Linux mount proxy system.

Before you begin

Run the uninstallation process as the root user. The root user profile must be sourced. If you use the **su** command to switch to root, use the **su -** command to source the root profile.

About this task

By default, when you uninstall the Linux mount proxy feature, the type of uninstallation that occurs is the same process as the type of original installation. To use a different uninstallation process, specify the correct parameter. For example, if you used a silent installation process, you can use the installation wizard to uninstall by specifying the **-i swing** parameter.

Procedure

To remove the Linux mount proxy feature, complete the following steps:

1. Change to the directory that contains the uninstallation program.
For example, issue the following command to change to the default location of the uninstallation program:

```
/opt/tivoli/tsm/_uninst/DPHyperV
```

2. Depending on the type of installation, use one of the following methods to uninstall the Linux mount proxy:

- To use the installation wizard to uninstall the Linux mount proxy, issue the following command:

```
./Uninstall_Data_Protection_for_Hyper-V -i swing
```

- To use the console to uninstall the Linux mount proxy, issue the following command:

```
./Uninstall_Data_Protection_for_Hyper-V -i console
```

- To silently uninstall the Linux mount proxy, issue the following command:

```
./Uninstall_Data_Protection_for_Hyper-V -i silent
```

Related tasks

[“Removing the file restore feature”](#) on page 38

If you no longer want to run file restore operations, you can remove the file restore feature by updating a configuration file. If you uninstall Data Protection for Microsoft Hyper-V, you must update the same configuration file to remove the file restore feature.

Removing the file restore feature

If you no longer want to run file restore operations, you can remove the file restore feature by updating a configuration file. If you uninstall Data Protection for Microsoft Hyper-V, you must update the same configuration file to remove the file restore feature.

About this task

To remove the services that are related to the file restore feature, you must update the `frConfig.props` file and remove the services that are related to the mount proxy nodes.

Procedure

Complete the following steps on the Hyper-V host or cluster:

1. Manually edit the
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\frConfig.props
file by changing the following option:

```
enable_filerestore=true
```

Change the option as shown:

```
enable_filerestore=false
```

2. Open the **Services** control panel in the Windows operating system and remove the services that are related to the mount proxy node.
Typically, the services are called TSM CAD - *hostname_HV_MP_platform* and TSM Agent - *hostname_HV_MP_platform*.
3. Remove the mount proxy nodes on the IBM Storage Protect server with the [REMOVE NODE](#) command.

Results

The file restore feature is removed from the Hyper-V host or cluster. You do not have to restart the IBM Storage Protect for Virtual Environments Derby Database or IBM Storage Protect Web Server services.

What to do next

If you ran file restore operations on a Linux guest virtual machine, you must uninstall Linux the mount proxy feature. For instructions, see [“Uninstalling the mount proxy feature on Linux systems”](#) on page 37.

Chapter 3. Configuring Data Protection for Microsoft Hyper-V

After successfully installing the IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V software, you must configure Data Protection for Microsoft Hyper-V before performing any backup and restore operations. If you want to run mount operations on guest VMs with the IBM Storage Protect recovery agent, you must also configure the recovery agent.

Tip: After you install Data Protection for Microsoft Hyper-V, the IBM License Metric Tool counts the data mover only if it is connected to an IBM Storage Protect server and is used for data operations. Subsequently, that data mover is always included in license calculations. Data movers that are not connected to a server and are not used for data operations are excluded from license calculations.

Configuring Data Protection for Microsoft Hyper-V with the wizard

You can use the configuration wizard for the initial configuration or to update the configuration of Data Protection for Microsoft Hyper-V for a stand-alone Hyper-V host or a cluster environment. You can also use the wizard to enable Data Protection for Microsoft Hyper-V for file restore operations.

Before you begin

- If you are upgrading from Data Protection for Microsoft Hyper-V 8.1.2 or earlier and have nodes that are already defined on the IBM Storage Protect server, complete the tasks in [“Renaming nodes on the IBM Storage Protect server”](#) on page 18.
- The Hyper-V host where Data Protection for Microsoft Hyper-V is installed must have network connectivity to the IBM Storage Protect server that is used to store virtual machine (VM) backups.
- To help improve performance, use at least a 10 Gb connection between the Hyper-V hosts and the IBM Storage Protect server.
- You must have the login credentials for the IBM Storage Protect server administrator account.
- You must connect to a secure IBM Storage Protect server that uses Secure Sockets Layer (SSL) communications. A security certificate is downloaded automatically when you are completing the configuration wizard.
- In a cluster environment, ensure that you install the Data Protection for Microsoft Hyper-V package on each host in the cluster. After you install the packages on all the hosts, run the installation wizard on one of the hosts in the cluster. The configuration wizard will connect to each host to complete the configuration.

Any node that does not have the Data Protection for Microsoft Hyper-V software installed is omitted from the cluster configuration, and does not affect the configuration of any other nodes where the software is installed. If you add a node to the cluster later, install Data Protection for Microsoft Hyper-V on that node and run the configuration wizard for that node (locally or from any other nodes in the cluster).

- On each host where Data Protection for Microsoft Hyper-V is installed, ensure that the HTTPS port that is used to communicate with Data Protection for Microsoft Hyper-V is open in the firewall. Unless specified otherwise, the default port number of 9081 is used.
- The configuration wizard determines the node names to use based on the host or cluster name. You can use the default node names or customize the node names by adding prefixes and suffixes. To customize the node names, you must complete the steps described in [“Customizing node names”](#) on page 20 before you run the configuration wizard.

About this task

To simplify the configuration, the configuration wizard automatically creates the nodes that are required for backup, restore, and optionally file restore operations. The configuration wizard also registers the nodes on the IBM Storage Protect server and configures the services on the local Windows host.

For more information about the types of nodes that are used for Data Protection for Microsoft Hyper-V, see [“How IBM Storage Protect nodes are used in Data Protection for Microsoft Hyper-V”](#) on page 6.

Procedure

To configure Data Protection for Microsoft Hyper-V, complete the following steps on the Hyper-V host. For a cluster environment, complete the following steps on any host in the cluster where Data Protection for Microsoft Hyper-V is installed.

1. Start the Data Protection for Microsoft Hyper-V Management Console by clicking **Start > IBM Storage Protect > DP for Hyper-V Management Console**.

Alternatively, issue the following command at the command prompt:

```
"C:\Program Files\IBM\StorageProtect\DPHyperV\DpHv.msc"
```

2. When prompted, log on to the Data Protection for Microsoft Hyper-V Management Console. Enter the same credentials that you use to log on to the Hyper-V host.

The account that you use must be a member of the local administrators group on the machine so that Hyper-V and cluster operations can be completed.

3. If you are configuring Data Protection for Microsoft Hyper-V Management Console for the first time, the configuration wizard opens automatically.

If you are changing the existing configuration in a stand-alone host environment, click a host in the navigation pane, and click **Configure** in the actions pane. In a cluster environment, click select a cluster node in the navigation pane, and click **Configure**.

4. Complete each page of the wizard and click **Next** to advance to the next page.

Page	Action
Before you Begin	Click Next to start the wizard.
Backup Server	<p>Enter information about the IBM Storage Protect server that is used to store VM backups.</p> <p>Backup server address The host name or IP address of the IBM Storage Protect server.</p> <p>Backup server SSL port Specify the port number for the server port that allows administrative connections by using the SSL protocol with TLS 1.2 enabled. The default port number is provided. Accept the default port number unless your server is configured to use a different port.</p> <p>Administrative credentials The user name and password of the IBM Storage Protect server administrator. The administrator must have system privilege, and be able to register client nodes on the server.</p>
Accept Certificate	<p>This window appears only if you are connecting to the IBM Storage Protect server for the first time, or if the existing security certificate is no longer valid. Click Accept to download and import the certificate automatically.</p> <p>If you are connecting to the version 8.1.1 or earlier V8 server, or version 7.1.7 or earlier server, and the download process fails, see “Configuring security settings to connect to IBM Storage Protect server version 8.1.1 or earlier or version 7.1.7 or earlier” on page 43.</p>

Page	Action
Cluster and Host Configuration	<p>The following options are available:</p> <p>Policy domain Select a policy domain from the list. The policy domain contains rules that determine how long VM backups are kept on the IBM Storage Protect server and how many versions of the VM backup are retained. The default policy domain is STANDARD.</p> <p>Target node name Displays the node name where VM backups are stored on the IBM Storage Protect server. For clusters, all VM backups are stored under the target node, regardless of which node in the cluster is running the backup.</p> <p>Node Definitions Displays the node definitions for the stand-alone host or hosts in the cluster. For information about the types of nodes, see Table 1 on page 7.</p> <p>Enable File Restore If you want to use the file restore web interface to restore individual files from a VM backup, check this box. When you select this check box, the mount proxy node pair for each host is automatically added to the list.</p> <p>This node pair represents the Linux and Windows proxy systems that access 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 for file restore operations.</p> <p>During the initial configuration, Enable File Restore is checked by default.</p> <p>File Restore Settings Click this button to enter the file restore administrator credentials.</p>
File Restore Settings	This window appears only if you enabled the file restore feature. Enter the file restore administrator credentials. The account must be a Windows domain user account with local administrative authority over all VMs.
Summary	Review the settings and click Next to complete the configuration.
Results	<p>The results of the configuration are displayed. If the configuration did not succeed, a list of errors is displayed. Correct the errors and run the configuration again.</p> <p>If the file restore feature is successfully configured, information about the host, Linux mount proxy, and file restore URLs are displayed in the file restore results table. You can click Copy to copy all the information to the clipboard.</p> <p>You must complete the file restore configuration by following the instructions in “Enabling the environment for file restore operations” on page 44.</p>

Results

Upon successful completion of the wizard, you can run backup and restore operations by using the command prompt, PowerShell cmdlets, or the Data Protection for Microsoft Hyper-V Management Console.

For your convenience, you can also open the file restore interface by clicking **File Restore** in the **Actions** pane.

What to do next

You can verify the configuration by running the Data Protection for Microsoft Hyper-V Management Console or the **Test-DpHvConfiguration** PowerShell cmdlet. For more information, see:

- [“Verifying the configuration of Data Protection for Microsoft Hyper-V” on page 71](#)
- [“Data Protection for Microsoft Hyper-V cmdlet examples” on page 138](#)

You can also specify the preferred host to log on to by using the **Set-DpHvMmcLoginPreferences** cmdlet. For more information, see [Chapter 8, “Protecting virtual machines by using Windows PowerShell cmdlets,” on page 133](#).

If you configured Data Protection for Microsoft Hyper-V with the file restore feature enabled and you ran the configuration wizard again after the initial configuration, the Linux mount proxy node password must be reset. To reset the password, use one of the following methods:

Method 1

On the Linux mount proxy, the IBM Storage Protect administrator runs the **dsmc** command and enters the IBM Storage Protect administrator user ID and password when prompted.

Method 2

Complete the following steps:

1. The IBM Storage Protect administrator resets the Linux mount proxy node password by running the **UPDATE NODE** server command on the IBM Storage Protect server console.
2. The Linux mount proxy node owner runs the **dsmc** command on the Linux mount proxy. When prompted, the owner enters the default Linux mount proxy node ID and new Linux mount proxy node password (obtained from the IBM Storage Protect server administrator).

Related tasks

[“Configuring non-default port numbers for Data Protection for Microsoft Hyper-V operations” on page 58](#)
If you do not want to use the default port numbers for the Data Protection for Microsoft Hyper-V web server or REST API services, you can configure different port numbers by using Windows PowerShell cmdlets.

Configuring security settings for Data Protection for Microsoft Hyper-V

The settings that are required to securely connect to the IBM Storage Protect server depend on the server version that you are connecting to.

About this task

IBM Storage Protect Version 8.1.2 or later and V7.1.8 servers provide an improved security protocol that uses Transport Layer Security (TLS) 1.2 to encrypt all communication between the server and clients. Data Protection for Microsoft Hyper-V and the server are automatically configured to communicate with each other by using the Secure Sockets Layer (SSL) protocol. Certificates are distributed automatically.

When you use the configuration wizard to configure Data Protection for Microsoft Hyper-V, you are prompted to accept the security certificate. No manual steps are required to obtain and import the certificate. For more information, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard” on page 39](#).

If you are connecting to earlier versions of IBM Storage Protect servers and the automatic download process of the certificate fails, you must manually download and import the security certificate before running the configuration wizard. For more information, see [“Configuring security settings to connect to IBM Storage Protect server version 8.1.1 or earlier or version 7.1.7 or earlier” on page 43](#).

Configuring security settings to connect to IBM Storage Protect server version 8.1.1 or earlier or version 7.1.7 or earlier

You can enable Data Protection for Microsoft Hyper-V to communicate with IBM Storage Protect server version 8.1.1 or earlier or version 7.1.7 or earlier with the Transport Layer Security (TLS) protocol.

About this task

If the server is configured to use SSL with TLS 1.2 enabled, a truststore with a certificate is created automatically by accepting the security certificate from the configuration wizard. However, if the automatic download process fails, you must manually create the truststore and run the configuration wizard again.

The following procedure uses the Java key and certificate management tool **keytool**.

This tool is in the C:\Program Files\Common Files\Tivoli\TSM\jvm80406\jre\bin directory. This location is subject to change based on the version of Java software that you are using.

Procedure

Complete the following steps on the stand-alone Hyper-V host. In a cluster environment, complete the following steps for each host in the cluster.

1. Obtain the necessary certificate from the IBM Storage Protect server administrator and download it to a location on your host, for example, the c:\cert directory.
2. From the command prompt, change to the truststore directory by issuing the following command:

```
cd
C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\truststores
```

If this folder does not exist, create it.

3. Import the certificate with the following command:

```
"C:\Program Files\Common Files\Tivoli\TSM\jvm80406\jre\bin\keytool.exe"
-importcert -alias my-cert -file "cert-filename" -keystore tsm-ve-
truststore.jks -storepass password
```

where:

-alias my-cert

The unique alias that identifies the certificate in the truststore.

-file "cert-filename"

The name of the file that contains the server self-signed certificate or the CA root certificate. For example, "C:\cert\cert256.arm".

-storepass password

The keystore password. Ensure that you remember this password for future use.

4. Start the Data Protection for Microsoft Hyper-V Management Console.

For instructions, see [“Starting the Data Protection for Microsoft Hyper-V Management Console” on page 61.](#)

5. Click **Configure** to open the configuration wizard.
6. On the **Backup Server** page, specify the port number in the **Backup server SSL port** field. This port is the server port that allows administrative connections by using SSL with TLS 1.2 enabled.
7. Complete the wizard.

Results

Upon successful completion of the wizard, you can run backup and restore operations by using the command prompt, PowerShell cmdlets, or the Data Protection for Microsoft Hyper-V Management Console.

Enabling the environment for file restore operations

When the file restore feature is enabled by an administrator, file owners can restore files with minimal assistance.

About this task

When you enable the file restore feature with the configuration wizard, the software that is needed for file restore operations is installed on the data mover node on a stand-alone Hyper-V host or on each host in a cluster.

In a cluster environment, the file restore software on each host in the cluster is independent of each other. In order for the file owner to be able to log on to the file restore interface, the host name and the virtual machine (VM) name that contains the file owner's data are required in the file restore URL.

Procedure

1. To start the configuration wizard, select a host or cluster from the navigation pane and click **Configure**.
2. Follow the instructions on each page of the wizard. For instructions, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard”](#) on page 39.
 - a) When you reach the **Cluster and Host Configuration** page, check the **Enable File Restore** check box.
 - b) If you are enabling file restore for the first time, you are prompted to enter the file restore administrator credentials. The administrator account must be a Windows domain user account with local administrative authority over all VMs.
3. Optional: If you plan to run file restore operations on Linux guest VMs, click **Copy** in the file restore results table in the **Results** page to copy the file restore URL and Linux mount proxy options to the clipboard. You can paste the mount proxy options to the `dsm.sys` file when you configure the Linux mount proxy.

You can also obtain this file restore information at any time after the configuration by clicking **Properties** in the **Actions** pane.

For more information about configuring the Linux mount proxy, see [“Configuring the Linux mount proxy for file restore operations”](#) on page 45.

4. Complete the configuration in the wizard.
5. Verify that you can access the file restore interface by selecting a VM from the **Results** pane and clicking **File Restore** in the **Actions** pane.
6. Construct the custom URL for each file owner based on the following template for the file restore URL:

```
https://<dphvhost>:9081/FileRestoreUI/login?  
vmName=<guestvm_name>&vmHost=<guestvm_host>&vmPlatform=<guestvm_platform>
```

where:

dphvhost

The Hyper-V host where you installed and configured Data Protection for Microsoft Hyper-V.

guestvm_name

The name of the guest VM that contains data for the file owner.

guestvm_host

The name of the VM host that is hosting the guest VM. The value for the *guestvm_host* can be the computer name, IP address, or DNS name.

guestvm_platform

The operating system of the guest VM. Specify one of the following values: **LINUX** or **WINDOWS**.

For example, if Data Protection for Microsoft Hyper-V is installed on a Hyper-V host called `Cluster1`, and the file owner's data is on a Windows guest VM called `MyVM-Win2k26` on VM host `HostB`, the file restore URL is as follows:

<https://Cluster1:9081/FileRestoreUI/login?vmName=MyVM-Win2k16&vmHost=HostB&vmPlatform=WINDOWS>

Fast path: You can also obtain the file restore URL by selecting a host, a VM, and clicking **File Restore** in the **Actions** pane. You can copy the URL address that is displayed in the web browser.

Tip: If you are using a non-default port number, replace port 9081 with the port that you configured. To show what port numbers that are being used, see [“Configuring non-default port numbers for Data Protection for Microsoft Hyper-V operations”](#) on page 58.

7. Distribute the file restore URL depending on the following scenarios:

- For the help desk model, the Hyper-V or file restore administrator sends a custom URL to each file owner.
- For the self-service model, the Hyper-V or file restore administrator sends instructions to file owners so they can construct their own file restore URLs. You can use the information about the URL from Step “6” on page 44 in your instructions to file owners.

Tip: VMs can fail over to different hosts in a cluster at any time. In this situation, you must send a new URL with the updated guest VM to the file owner, or the file owner must contact you to determine which host is hosting the VM.

Results

File owners are able to log in to the file restore interface to restore individual files and folders.

Configuring the Linux mount proxy for file restore operations

To prepare a Linux guest virtual machine (VM) for file restore operations, you must configure the Linux mount proxy.

Before you begin

Ensure that you complete the following tasks:

1. Run the Data Protection for Microsoft Hyper-V configuration wizard on the Hyper-V host or cluster and enable the file restore feature. For instructions, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard”](#) on page 39.
2. Install the mount proxy on a Linux system. For instructions, see [“Installing the mount proxy feature on Linux systems”](#) on page 34.

About this task

When you enable the file store feature with the configuration wizard, the mount proxy node pair for a stand-alone host or for each host in a cluster is automatically registered with the IBM Storage Protect server, along with the definition of proxy relationships. To enable the guest VM for file restore operations, you must configure the Linux mount proxy by adding the Linux mount proxy options that are provided by the configuration wizard to the `dsm.sys` file.

The following procedure sets up the mount proxy node by updating the mount proxy node options and verifying connectivity to the IBM Storage Protect server.

Procedure

Complete the following steps on the Linux mount proxy system:

1. If the client-user options file (`dsm.opt`) is not in the installation directory (`opt/tivoli/tsm/client/ba/bin`), create the file with a text editor.
2. Open the `dsm.opt` file with a text editor and add the following statement to the file:

```
servername MPNODE_hostname_HV_MP_LNX
```


where *hostname* is the name of the Windows Hyper-V host.

Ensure that this statement is the only statement in the file. Save your updates and close the file.

3. Open the `dsm.sys` file with a text editor. Copy the mount proxy options from the **Linux Mount Proxy Options** window of the configuration wizard and paste them into the file.

For example, paste the following stanza to the `dsm.sys` file:

```
SERVERNAME      MPNODE_<hostname>_HV_MP_LNX
NODename        <hostname>_HV_MP_LNX
PASSWORDAccess  generate
TCPServeraddress <backup_server_address>
TCPPort         1500
HTTPPort        1581    ** Must be unique for each node
COMMMethod      tcpip
ERRORLOGName    dsmerror.<hostname>_HV_MP_LNX.log
```

where *hostname* is the name of the Hyper-V host and *backup_server_address* is the host name or IP address of the IBM Storage Protect server where VMs are backed up.

Save your changes and close the `dsm.sys` file.

4. Start a command-line session on the mount proxy system with the `-asnodename` and `-optfile` command-line parameters:

```
dsmc -asnodename=hyperv_target_node -optfile=dsm.opt
```

where *hyperv_target_node* is the Hyper-V node name under which your VM backups are stored. The Hyper-V target node has the following naming convention:

- For a stand-alone host environment: *hostname_HV_TGT*
- For a cluster environment: *clustername_HV_TGT*

During the initial sign-on, you are prompted for a user ID and password. Enter your IBM Storage Protect server administrator ID and password.

After the initial sign-on, a new password is generated and stored so that you will not be prompted for the password again.

To ensure that you are not prompted for the password, run the **dsmc** command again. If you are prompted for the password, ensure that the `passwordaccess generate` option is set in the `dsm.sys` file, and repeat Step “4” on page 46 again.

5. Verify the connection to the IBM Storage Protect server by issuing the following command:

```
dsmc query session
```

This command shows information about your session, including the current node name, the session start time, server information, and server connection information.

6. Set up the client acceptor service (CAD) by taking the following actions:

- a) Set the following environment variable in the `/etc/init.d/dsmcad` file:

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

- b) The installation program creates a startup script for the client acceptor (`dsmcad`) in the `/etc/init.d` directory. The client acceptor must be started before it can manage scheduler tasks.

Ensure that you are logged in with the root user ID, and then use the following command to start the client acceptor:

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

```
service dsmcad start
```

To enable the client acceptor to start automatically after a system restart, add the service as follows, at a shell prompt:


```
# chkconfig --add dsmscad
```

What to do next

Note the messages on the status of various aspects of your file restore operating environment.

Follow the guidelines for required tasks such as disabling multipathing and Logical Volume Manager filtering.

Modifying options for file restore operations

To allow administrators to configure and control file restore operations, modify the options in the `frConfig.props` file.

About this task

Complete these steps on the system where the Data Protection for Microsoft Hyper-V Management Console is installed.

Procedure

1. Go to the directory where the `frConfig.props` file is located.
For example, open a command prompt and issue the following command:

```
cd C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI
```

2. Open the `frConfig.props` file with a text editor in administrator mode and modify the options as needed.
To determine which options to modify, see [“Options for file restore operations” on page 48](#).
3. Save your changes and close the `frConfig.props` file.

Results

Modified options are applied to the IBM Storage Protect file restore interface.

Options for file restore operations

The options in the `frConfig.props` file control configuration, support, and restore processing for file restore operations.

backup_info_duration_hours=num_hrs

Specify the amount of time, in hours, that information about recent backup activity is retained in the local Data Protection for Microsoft Hyper-V Derby database. The maximum value is 14 days (336 hours). The default value is one week (168 hours).

enable_contact_info=false | true

Specify whether to provide administrator contact information that file owners can use to obtain support in the IBM Storage Protect file restore interface.

false

File owners do not receive administrator contact information. This value is the default.

true

File owners receive administrator contact information.

If you specify **enable_contact_info=true**, you must provide information in the **contact_info** option.

enable_filerestore=false | true

Specify whether file owners can restore their files from a virtual machine with the IBM Storage Protect file restore interface.

false

File owners cannot restore their files with the IBM Storage Protect file restore interface. This value is the default.

true

File owners can restore their files with the IBM Storage Protect file restore interface.

maximum_mount_points=num_mount_points

Specify the maximum number of simultaneous recovery points that are available to the user account. The minimum value is 1 recovery point. The maximum value is 256 mount points. The default value is 2 mount points.

Tip: To prevent a virtual machine from being mounted multiple times for simultaneous restore operations, set this option with a low value.

mount_session_timeout_minutes=num_mins

Specify the amount of time, in minutes, that a restore and the mounted recovery point can be idle before the session is canceled. A cancellation unmounts the recovery point. The maximum value is 8 hours (480 minutes). The default value is 30 minutes.

Tip: To prevent the session from being canceled unexpectedly, increase the number of minutes.

restore_info_duration_hours=num_hrs

Specify the amount of time, in hours, that information about recent restore activity is retained for the IBM Storage Protect file restore interface. Use the restore activity window to view error information and recently completed tasks. This information provides a way to locate recently restored files. The maximum value is 14 days (336 hours). The default value is one week (168 hours).

contact_info=administrator information

Provide administrator contact information that file owners can use to obtain support. Contact information displays in the IBM Storage Protect file restore interface in the following locations:

- Login window
- The **About** pane in the help menu
- The support information link in interface messages

The Data Protection for Microsoft Hyper-V Management Console wizard can overwrite the **enable_filerestore** option, but only to the **true** value. You must manually set the option to **false** if you want to disable the file restore feature.

Configuring Data Protection for Microsoft Hyper-V log activity

To allow administrators to configure and control how content is formatted and logged for Data Protection for Microsoft Hyper-V Management Console and file restore operations, modify the options in the FRLog.config file.

Before you begin

The FRLog.config file is generated the first time that the Data Protection for Microsoft Hyper-V Management Console or file restore interface is accessed.

About this task

Complete these steps on the system where the Data Protection for Microsoft Hyper-V Management Console is installed.

Procedure

1. Go to the directory where the FRLog.config file is located.
Open a command prompt and issue the following command:

```
cd C:\IBM\StorageProtect\webserver\usr\servers\veProfile\frGUI
```

2. Open the FRLog.config file with a text editor in administrator mode and modify the options as needed.

To determine which options to modify, see [“Data Protection for Microsoft Hyper-V log activity options”](#) on page 50.

3. Save your changes and close the `FRLog.config` file.
4. Restart the GUI web server:
 - a) Click **Start > Control Panel > System and Security > Administrative Tools > Services**.
 - b) Right-click **IBM Storage Protect for Virtual Environments Web Server** and click **Restart**.

Results

Settings are applied to the content and format of logging information for Data Protection for Microsoft Hyper-V Management Console and file restore operations.

Data Protection for Microsoft Hyper-V log activity options

The `FRLog.config` options control the content and format of logging information for Data Protection for Microsoft Hyper-V Management Console and file restore operations.

The following options log information for Data Protection for Microsoft Hyper-V Management Console and file restore tasks in the `fr_gui.log` file:

MAX_LOG_FILES=number

Specify the maximum number of `fr_gui.log` files to retain. The default value is 8.

MAX_LOG_FILE_SIZE=number

Specify the maximum size of the `fr_gui.log` file in KBs. The default value is 8192 KB.

The following options log information for Data Protection for Microsoft Hyper-V Management Console and file restore services in the `fr_api.log` file. These services are internal API services that are related to Data Protection for Microsoft Hyper-V Management Console and file restore activity:

API_MAX_LOG_FILES=number

Specify the maximum number of `fr_api.log` files to retain. The default value is 8.

API_MAX_LOG_FILE_SIZE=number

Specify the maximum size of the `fr_api.log` file in KBs. The default value is 8192 KB.

API_LOG_FILE_NAME=API_log_file_name

Specify the name of the API log file. The default value is `fr_api.log`.

API_LOG_FILE_LOCATION=API_log_file_name

Specify the location of the API log file. The location must be specified with a forward slash (/).

The default location is `install_directory/IBM/StorageProtect/webserver/usr/servers/veProfile/logs`.

FR.API.LOG=ON | OFF

Specify whether to enable logging for Data Protection for Microsoft Hyper-V Management Console and file restore services.

- To enable logging, specify ON. The default value is ON.
- To disable logging, specify OFF.

Configuring the IBM Storage Protect recovery agent GUI

You must set up the IBM Storage Protect recovery agent GUI for mount and file restore operations.

Before you begin

These configuration tasks must be completed before you use the IBM Storage Protect recovery agent GUI.

Procedure

1. Log on to the system where you want to restore files. The IBM Storage Protect recovery agent must be installed on the system.

2. Click **Select IBM Storage Protect server** in the IBM Storage Protect recovery agent GUI to connect to the IBM Storage Protect server.

Specify the following options:

Server address

Enter the IP address or host name of the IBM Storage Protect server.

Server port

Enter the port number that is used for TCP/IP communication with the server. The default port number is 1500.

Node access method:

Asnodename

Select this option to use a proxy node to access the virtual machine backups that are in the target node. The proxy node is a node that is granted "proxy" authority to perform operations on behalf of the target node.

Typically, you use the `grant proxynode` command to create the proxy relationship between two existing nodes.

If you select this option, complete the following steps:

- a. Enter the name of the target node (the node where the virtual machine backups are located) in the **Target Node** field.
- b. Enter the name of the proxy node in the **Authentication node** field.
- c. Enter the password for the proxy node in the **Password** field.
- d. Click **OK** to save these settings and exit the IBM Storage Protect page.

When you use this method, the IBM Storage Protect recovery agent user knows only the proxy node password, and the target node password is protected.

Fromnode

Select this option to use a node with access limited only to the snapshot data of specific virtual machines in the target node.

Typically, this node is given access from the target node that owns the virtual machine backups by using the `set access backup` command:

```
set access backup -TYPE=VM vmdisplayname mountnodename
```

For example, this command gives the node named `myMountNode` the authority to restore files from the virtual machine named `myTestVM`:

```
set access backup -TYPE=VM myTestVM myMountNode
```

If you select this option, complete the following steps:

- a. Enter the name of the target node (the node where the virtual machine backups are located) in the **Target Node** field.
- b. Enter the name of the node that is given limited access in the **Authentication node** field.
- c. Enter the password for the node that is given limited access in the **Password** field.
- d. Click **OK** to save these settings and exit the IBM Storage Protect page.

When you use this method, you can see a complete list of backed-up virtual machines. However, you can restore only those virtual machine backups to which the node was granted access. In addition, the snapshot data is not protected from expiration on the server.

Direct

Select this option to authenticate directly to the target node (the node where the virtual machine backups are located).

If you select this option, complete the following steps:

- a. Enter the name of the target node (the node where the virtual machine backups are located) in the **Authentication node** field.
- b. Enter the password for the target node in the **Password** field.
- c. Click **OK** to save these settings and exit the IBM Storage Protectpage.

Use Password access generate

When this option is selected and the password field is empty, the IBM Storage Protect recovery agent authenticates with an existing password that is stored in the password store. If not selected, you must manually enter the password.

To use this option, you must first manually set an initial password for the node to which the option applies. You must specify the initial password when you connect to the IBM Storage Protect node for the first time by entering the password in the **Password** field and selecting the **Use Password access generate** check box.

However, when you use the local data mover node as the **Authentication node**, the password might already be stored in the password store. As a result, select the **Use Password access generate** check box and do not enter a password.

For more information about the password store, see [Secure password storage](#).

The IBM Storage Protect recovery agent queries the specified server for a list of protected virtual machines, and shows the list.

3. Set the following mount, backup, and restore options by clicking **Settings**:

Virtual Volume write cache

The IBM Storage Protect recovery agent that is running on the backup proxy host saves data changes on a virtual volume in the write cache. By default, the write cache is enabled and the maximum cache size is 90% of the available space for the selected folder. To prevent the system volume from becoming full, change the write cache to a path on a volume other than the system volume.

Folder for temporary files

Specify the path where data changes are saved. The write cache must be on a local drive and cannot be set to a path on a shared folder.

Cache size

Specify the size of the write cache. The maximum allowed cache size is 90% of the available space for the selected folder.

Restriction: To prevent any interruption during restore processing, exclude the write cache path from all antivirus software protection settings.

Data Access

Specify the type of data to be accessed. If you are using an offline device (such as tape or virtual tape library), you must specify the applicable data type.

Storage type

Specify one of the following storage devices from which to mount the snapshot:

Disk/File

The snapshot is mounted from a disk or file. This device is the default.

Tape

The snapshot is mounted from a tape storage pool. When this option is selected, it is not possible to mount multiple snapshots.

VTL

The snapshot is mounted from an offline virtual tape library. Concurrent mount sessions on the same virtual tape library are supported.

Requirement: When the storage type is changed, you must restart the service for the changes to take effect.

Disable expiration protection

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.

- To protect the snapshot from expiration, do not select this option. This option is cleared by default. The snapshot on the IBM Storage Protect server is locked and the snapshot is protected from expiration during the mount operation.
- To disable expiration protection, select this option. The snapshot on the IBM Storage Protect 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 selecting this option.

Read Ahead size (in 16-KB blocks)

Specify the number of extra data blocks that are retrieved from the storage device after a read request is sent to a single block. The default values are as follows:

- Disk or file: 64
- Tape: 1024
- VTL: 64

The maximum value for any device is 1024.

Read Ahead cache size (in blocks)

Specify the size of the cache where the extra data blocks are stored. The default values are as follows:

- Disk or file: 10000
- Tape: 75000
- VTL: 10000

Since each snapshot has its own cache, make sure to plan how many snapshots are mounted or restored simultaneously. The cumulative cache size cannot exceed 75000 blocks.

Driver timeout (seconds)

This value specifies the amount of time to process data requests from the file system driver. If processing is not completed in time, the request is canceled and an error is returned to the file system driver. Consider increasing this value when you experience timeouts. For example, timeouts might occur when the network is slow, the storage device is busy, or multiple mount sessions are being processed. The default values are as follows:

- Disk or file: 60
- Tape: 180
- VTL: 60

Click **OK** to save your changes and exit the **Settings**.

4. Verify that each IBM Storage Protect server node (that was specified with the `Asnodename` and `Fromnode` options) allows backups to be deleted.

The IBM Storage Protect recovery agent creates unused temporary objects during operations. The `BACKDELeTe=Yes` server option allows these objects to be removed so that they do not accumulate in the node.

- a) Log on to the IBM Storage Protect server and start an administrative client session in command-line mode:

```
dsmadm -id=admin -password=admin -dataonly=yes
```

- b) Enter the following command:

```
Query Node <nodename> Format=Detailed
```

Make sure the command output for each node includes the following statement:

```
Backup Delete Allowed?: Yes
```

If this statement is not included, update each node with this command:

```
UPDate Node <nodename> BACKDELeTe=Yes
```

Run the Query Node command again for each node to verify that each node allows backups to be deleted.

Enabling secure communication from the recovery agent to the IBM Storage Protect server

If the IBM Storage Protect server is configured to use the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocol, you can enable the recovery agent to communicate with the server by using the protocol.

Before you begin

Consider the following requirements before you begin configuration for secure communication to the server:

- Each server that is enabled for SSL must have a unique certificate. The certificate can be one of the following types:
 - A certificate that is self-signed by the server.
 - A certificate that is issued by a third-party certificate authority (CA) certificate. The CA certificate can be from a company such as Symantec or Thawte, or an internal certificate that is maintained within your company.
- For performance reasons, use SSL or TLS only for sessions where security is required. Consider adding more processor resources on the server system to manage the increased requirements.
- For a client to connect to a server that is using TLS version 1.2, the certificate signature algorithm must be Secure Hash Algorithm 1 (SHA-1) or later. If you are using a self-signed certificate to a server that is using TLS V1.2, you must use the `cert256.arm` certificate. Your IBM Storage Protect administrator might need to change the default certificate on the server.
- To disable security protocols that are less secure than TLS 1.2, add the **SSLDISABLELEGACYt1s yes** option to the `C:\windows\system32\fb.opt` or `C:\Windows\SysWOW64\fb.opt` file. TLS 1.2 or later helps to prevent attacks by malicious programs.

Enabling secure communication by using an IBM Storage Protect server self-signed certificate

If the IBM Storage Protect server is using a self-signed certificate, you must obtain a copy of that certificate from the server administrator and configure the recovery agent to communicate with the server by using the SSL or TLS protocol.

About this task

Each server generates its own certificate. Version 6.3 and later servers generate files that are named `cert256.arm` if the server is using TLS 1.2 or later or `cert.arm` if the server is using an earlier version of SSL or TLS. Server versions earlier than V6.3 generate files that are named `cert.arm` regardless of the protocol. You must choose the certificate that is set as the default on the server.

The certificate file is stored on the server workstation in the server instance directory. For example, `C:\IBM\tivoli\tsm\server\bin\cert256.arm`. If the certificate file does not exist, the certificate file is created when you restart the server with these options set.

Procedure

To enable SSL or TLS communication from the recovery agent to the server by using a self-signed certificate:

1. Append the GSKit binary path and library path to the PATH environment variable on the client.

For example:

```
set PATH=C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\bin\;  
C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\lib64;%PATH%
```

2. If you are configuring SSL or TLS on the client for the first time, you must create the client local key database `dsmcert.kdb`.

From the `C:\Windows\SysWOW64` directory, run the **`gsk8capicmd_64`** command as shown in the following example:

```
gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw password -stash
```

The password that you provide is used to encrypt the key database. The password is automatically stored encrypted in the stash file (`dsmcert.sth`). The stash file is used by the client to retrieve the key database password.

3. Obtain the server self-signed certificate.
4. Import the certificate in to the `dsmcert.kdb` database. You must import the certificate for each client in to the `dsmcert.kdb`.

From the `C:\Windows\SysWOW64` directory, run the **`gsk8capicmd_64`** command as shown in the following example:

```
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -label "Server server_name self-signed  
key"  
-file path_to_certificate -format ascii -trust enable
```

Multiple server certificates can be added to the `dsmcert.kdb` database so that the client can connect to different servers. Different certificates must have different labels. Use meaningful names for the labels.

Important: For a disaster recovery of the server, if the certificate has been lost, the server automatically generates a new certificate. Each client must then import the new certificate.

5. After the server certificate is added to the `dsmcert.kdb` database, add the `ssl yes` option to the `C:\Windows\SysWOW64\fb.opt` file and update the value of the `tcpport` option.

Important:

The server is normally set up for SSL and TLS connections on a different port than non-SSL and TLS connections. Do not specify a non-SSL or TLS port number for the `tcpport` value. If the value of `tcpport` is incorrect, the recovery agent cannot connect to the server.

You cannot connect to a non-SSL or TLS port with a recovery agent that is enabled for SSL or TLS or connect a SSL or TLS port to a recovery agent that is not enabled for SSL or TLS.

6. Set the correct SSL or TLS ports in the following recovery agent configuration files:

- `C:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\RecoveryAgent.conf`
- `C:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\RecoveryAgentDMNodes.conf`

Enabling secure communication by using a third-party certificate

If the IBM Storage Protect server is using a third-party certificate authority (CA), you must obtain the CA root certificate.

About this task

If the certificate was issued by a CA such as Symantec or Thawte, the client is ready for SSL or TLS and you can skip the following configuration steps. For a list of preinstalled CA root certificates, search for **Certificate Authorities root certificates** on the IBM IBM Documentation.

If the certificate was not issued by a preinstalled root certificate or is an internal CA certificate that is maintained within your company, you must configure the recovery agent to communicate with the server by using the SSL or TLS protocol.

Procedure

To enable SSL or TLS communication from the recovery agent to the server by using a CA certificate:

1. Append the GSKit binary path and library path to the `PATH` environment variable.

For example:

```
set PATH=C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\bin\;  
C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\lib64;%PATH%
```

2. If you are configuring SSL or TLS on the client for the first time, you must create the client local key database `dsmcert.kdb`.

For clients, from the `C:\Windows\SysWOW64` directory, run the **`gsk8capicmd_64`** command as shown in the following example:

```
gsk8capicmd_64 -keydb -create -populate -db dsmcert.kdb -pw password -stash
```

The password that you provide is used to encrypt the key database. The password is automatically stored encrypted in the stash file (`dsmcert.sth`). The stash file is used by the client to retrieve the key database password.

3. Obtain the CA certificate.

4. Import the certificate in to the `dsmcert.kdb` database. You must import the certificate for each client in to the `dsmcert.kdb`.

For clients, from the `C:\Windows\SysWOW64` directory, run the **`gsk8capicmd_64`** command as shown in the following example:

```
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed -label "XYZ Certificate Authority"  
-file path_to_CA_root_certificate -format ascii -trust enable
```

Multiple server certificates can be added to the `dsmcert.kdb` database so that the client can connect to different servers. Different certificates must have different labels. Use meaningful names for the labels.

Important: For a disaster recovery of the server, if the certificate has been lost, the server automatically generates a new certificate. Each client must import the new certificate.

5. After the server certificate is added to the dsmcert.kdb database, add the `ssl yes` option to the `C:\Windows\SysWOW64\fb.opt` file and update the value of the `tcpport` option.

Important:

The server is normally set up for SSL and TLS connections on a different port than non-SSL and TLS connections. Do not specify a non-SSL or TLS port number for the `tcpport` value. If the value of `tcpport` is incorrect, the recovery agent cannot connect to the server.

You cannot connect to a non-SSL or TLS port with a recovery agent that is enabled for SSL or TLS or connect a SSL or TLS port to a recovery agent that is not enabled for SSL or TLS.

6. Set the correct SSL or TLS ports in the following recovery agent configuration files:
 - `C:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\RecoveryAgent.conf`
 - `C:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\RecoveryAgentDMNodes.conf`

Manually configuring an iSCSI device

You must configure the Windows system that is used during an iSCSI mount operation. The snapshot is mounted from IBM Storage Protect server storage.

Before you begin

Review the following iSCSI requirements before you proceed with this task:

- During an iSCSI mount, an iSCSI target is created on the IBM Storage Protect recovery agent system. You can connect to the iSCSI target from any system to create a volume that contains the backup data. Also, you can then mount this volume from another system.
- iSCSI initiator is required on any system that must connect to the iSCSI target.
- Make sure that an iSCSI initiator is installed on the system where the data is to be restored.
- Microsoft iSCSI Initiator is not required on the IBM Storage Protect recovery agent system.

Review the following disk and volume requirements before you proceed with this task:

- If a volume spans several disks, you must mount all the required disks. When mirrored volumes are used, mount only one of the mirrored disks. Mounting one disk prevents a time-consuming synchronization operation.
- If multiple dynamic disks were used on the backup system, these disks are assigned to the same group. As a result, Windows Disk Manager might consider some disks as missing and issue an error message when you mount only one disk. Ignore this message. The data on the backed up disk is still accessible, unless some of the data is on the other disk. This issue can be solved by mounting all the dynamic disks.

About this task

Complete these tasks to configure the Windows system that is used during an iSCSI mount operation:

Procedure

1. On the IBM Storage Protect recovery agent system, open port 3260 in the LAN firewall and the Windows client firewall.

Record the iSCSI initiator name on the system where data is to be restored.

The iSCSI initiator name is shown in the iSCSI initiator configuration window of the Control Panel. For example:

```
iqn.1991-05.com.microsoft:hostname
```

2. Complete these tasks on the system where the IBM Storage Protect recovery agent (or iSCSI target) is installed:

- a) Start the IBM Storage Protect recovery agent GUI. Complete the **Select IBM Storage Protect server** and **Select snapshot** dialogs and click **Mount**.
 - b) In the **Choose mount destination** dialog, select **Mount an iSCSI target**.
 - c) Create a target name. Make sure that it is unique and that you can identify it from the system that runs the iSCSI initiator. For example:

iscsi-mount-tsm4ve
 - d) Enter the iSCSI Initiator name that was recorded in Step 1 and click **OK**.
 - e) Verify that the volume you just mounted is displayed in the **Mounted Volumes** field.
3. Locate and start the iSCSI Initiator program on the initiator system that was selected in Step 1:
 - a) Connect to the iSCSI target:
 - i) In the Targets tab, enter the TCP/IP address of the IBM Storage Protect recovery agent (iSCSI target) used in Step 2 in the **Target:** dialog. Click **Quick Connect**.
 - ii) The **Quick Connect** dialog shows a target that matches the target name that was specified in Step 2c. If it is not already connected, select this target and click **Connect**.
 - b) On the initiator system, go to **Control Panel > Administrative Tools > Computer Management > Storage > Disk Management**.
 - i) If the mounted iSCSI target is listed as **Type=Foreign**, right-click **Foreign Disk** and select **Import Foreign Disks**. The **Foreign Disk Group** is selected. Click **OK**.
 - ii) The next screen shows the type, condition, and size of the Foreign Disk. Click **OK** and wait for the disk to be imported.
 - iii) When the disk import completes, press **F5** (refresh). The mounted iSCSI snapshot is visible and contains an assigned drive letter. If drive letters are not automatically assigned, right-click the required partition and select **Change Drive Letters or Paths**. Click **Add** and select a drive letter.
 4. Open Windows Explorer (or other utility) and browse the mounted snapshot for a file restore operation.
 5. After the file is restored, complete these tasks:
 - a) Disconnect each iSCSI target by using the **iSCSI Initiator Properties** dialog.
 - b) Dismount the volume from Step 2 by selecting the volume in the IBM Storage Protect recovery agent GUI and clicking **Dismount**.

Advanced configuration

Use advanced configuration tasks to further customize the configuration of Data Protection for Microsoft Hyper-V.

Configuring non-default port numbers for Data Protection for Microsoft Hyper-V operations

If you do not want to use the default port numbers for the Data Protection for Microsoft Hyper-V web server or REST API services, you can configure different port numbers by using Windows PowerShell cmdlets.

About this task

The default port number that is assigned to the web server provides services to the Data Protection for Microsoft Hyper-V Management Console, file restore interface, and PowerShell cmdlets.

Complete the steps in the following procedure to change the port number.

Procedure

1. Start PowerShell by following the instructions in [“Preparing to use PowerShell cmdlets with Data Protection for Microsoft Hyper-V”](#) on page 133.

2. Optional: Show what port numbers that are being used by running the following cmdlets:

- To show the web server port, use the **Show-DpHvHttpsPort** cmdlet.
- To show the REST API port, use the **Show-DpHvMmcLoginPreferences** cmdlet.

This cmdlet shows the login preferences, including the REST API port number, for the Data Protection for Microsoft Hyper-V Management Console. The preferences are created when the management console is run the first time. If you run this cmdlet before the management console is ever run, no information is returned.

3. To change the default port numbers, use the following cmdlets:

- To change the web server port, use the **Set-DpHvHttpsPort** cmdlet. For example, to change the web server port number to 9082, use the following cmdlet:

```
Set-DpHvHttpsPort -httpsPort 9082
```

All hosts in a cluster must use the same HTTPS port.

- To change the REST API port, use the **Set-DpHvMmcLoginPreferences** cmdlet. For example, to change the REST API port number to 9082, use the following cmdlet:

```
Set-DpHvMmcLoginPreferences -RestApiPort 9082
```

Tip: For more examples, use the **Get-Help cmdlet_name** command.

Tuning scheduled VM backups for Windows Server 2012 and 2012 R2 clusters

Beginning with Data Protection for Microsoft Hyper-V version 8.1.2, you can back up more virtual machines (VMs) in parallel and across nodes in a cluster. A cluster node backup operation always retries the snapshot on volumes with snapshots that failed with a recoverable condition. You can also tune the number of VMs in a snapshot to reduce the workload of a snapshot for the Hyper-V host.

You can use the following options to tune how snapshots are taken during the backup:

- Use the `vmmaxparallel` option to control how many VMs are sent in parallel to the IBM Storage Protect server. The setting for this option has the most notable impact on performance.
- Use the `vmmaxpersnapshot` option to control how many VMs can be included in each snapshot that is created during the backup operation.

Before you back up a cluster, review and tune the values for these two options for the environment.

Use the following general approach to tune your cluster backup operations:

1. Plan to use an appropriately sized and configured IBM Storage Protect server that uses container pools. For information about how to size the server, see [IBM Storage Protect Blueprints](#).
2. As a starting point, use the default values for the `vmmaxpersnapshot` and `vmmaxparallel` options.
3. Run the backup schedule and note the results, such as whether backups completed within the schedule window or whether too many snapshot retries occurred.
4. Adjust the value for the `vmmaxparallel` option to work in your environment. For example, set the value to 10.
5. Adjust the value of `vmmaxpersnapshot` to a value that minimizes the number of retries that occur. The retries are reported in the backup statistics.

When you choose a smaller number of VMs per snapshot, you increase the number of snapshots that are needed to complete a backup operation. This increase in snapshots can lead to delays during cluster backup operations of VMs on CSVs. The delay occurs because only one snapshot can be

created at a time, and backup operations of other nodes in the schedule are delayed during snapshot creation. By increasing the number of VMs in a snapshot, you can reduce the number of snapshots that are taken for a backup operation.

To determine the number of VMs to include in a snapshot, consider the following factors:

- A snapshot with more VMs takes longer to complete and increases the load on the system. A larger number of VMs means that the snapshot persists longer, which can affect system performance.
- The `vmmaxpersnapshot` and `vmmaxparallel` options work together to determine how many snapshots are taken in a backup operation. The `vmmaxparallel` option specifies how many VMs can be backed up simultaneously. Data Protection for Microsoft Hyper-V takes as many snapshots as needed to meet the `vmmaxparallel` setting.

VMs are sorted and selected based on the volumes that are needed to create the snapshot for the VMs. A snapshot is created for a set of VMs that share a set of volumes. Thus, the number of snapshots varies depending upon the volumes that are used by the VMs. The number of VMs per snapshot never exceeds the value for the `vmmaxpersnapshot` option.

The following table shows examples of how many VMs can be processed per snapshot with various `vmmaxpersnapshot` and `vmmaxparallel` settings. In these examples, assume that all the VMs are on the same volume.

<i>Table 7. Number of snapshots and VMs (on the same volume) processed with the <code>vmmaxpersnapshot</code> and <code>vmmaxparallel</code> settings</i>		
vmmaxpersnapsho t setting	vmmaxparallel setting	Number of snapshots created
10	20	Two snapshots are created with 10 VMs each. When the number of VMs being processed is less than the <code>vmmaxparallel</code> setting, another snapshot is taken.
20	20	One snapshot is created containing 20 VMs.
20	10	One snapshot is created containing 20 VMs, and 10 VMs are backed up due to the <code>vmmaxparallel</code> setting during the first run. The remaining 10 VMs are backed up during the second run (a second snapshot is not needed).

You can also use the `vmmaxsnapshotretry` option to specify the maximum number of times to retry a snapshot operation of a VM if the initial snapshot fails with a recoverable condition.

Related concepts

[“Limitations on Hyper-V backup operations” on page 10](#)

Before you start a Hyper-V backup operation, review the limitations. Some limitations apply to all Hyper-V backup operations, while others apply only to Hyper-V backups on Windows Server 2012 or 2012 R2 or Windows Server 2016 environments.

Related reference

[“Vmmaxpersnapshot” on page 195](#)

Use the `vmmaxpersnapshot` option to specify the maximum number of virtual machines (VMs) to include in a Hyper-V snapshot. The VMs in the snapshot are backed up to the IBM Storage Protect server.

[“Vmmaxsnapshotretry” on page 203](#)

Use the `vmmaxsnapshotretry` option to specify the maximum number of times to retry a snapshot operation of a virtual machine (VM) if the initial snapshot fails with a recoverable condition.

[“Vmmaxparallel” on page 194](#)

The `vmmaxparallel` option is used to configure parallel backups of several virtual machines (VMs) by using a single instance of the data mover. This option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at a time.

Chapter 4. Managing data with the Data Protection for Microsoft Hyper-V Management Console

The Data Protection for Microsoft Hyper-V Management Console provides a single environment to help you manage the daily operations of Data Protection for Microsoft Hyper-V.

You can use the Data Protection for Microsoft Hyper-V Management Console to start ad hoc backup and restore operations and to view the most recent backup information for all virtual machines (VMs) that are in a Hyper-V host or cluster.

This information includes the identification of VMs that are at risk of being unprotected because the VM has never been backed up or because a backup did not occur in the time interval that is set in the at-risk policy. The at-risk policy only applies to VMs that have been previously backed up.

Tip: You can also use the configuration wizard to initially configure or update the configuration of Data Protection for Microsoft Hyper-V. For more information, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard”](#) on page 39.

Starting the Data Protection for Microsoft Hyper-V Management Console

To manage day-to-day operations for Data Protection for Microsoft Hyper-V, start the Data Protection for Microsoft Hyper-V Management Console and enter your logon credentials.

Procedure

1. Start the Data Protection for Microsoft Hyper-V Management Console by clicking **Start > IBM Storage Protect > DP for Hyper-V Management Console**.

Alternatively, issue the following command at the command prompt:

```
"C:\Program Files\IBM\StorageProtect\DPHyperV\DpHv.msc"
```

2. When prompted, log on to the Data Protection for Microsoft Hyper-V Management Console. Enter the same credentials that you use to log on to the Hyper-V host.

The account that you use must be a member of the local administrators group on the computer so that Hyper-V and cluster operations can be completed.

Tip: If you did not configure Data Protection for Microsoft Hyper-V or if the configuration is incomplete, the configuration wizard appears automatically. For more information, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard”](#) on page 39.

3. If the security certificate that is associated with the host your are connecting to is not recognized or is not on the server where the Data Protection for Microsoft Hyper-V Management Console is installed, you are prompted to install a new certificate.

Complete the following steps for a stand-alone host or for each host in a cluster:

- a) In the **Securing connection to <host name>** window, click **View certificate**.

If you select any other options, such as **Yes** to ignore the certificate warning for the current session, **No** to stop the connection, or **Don't ask me again for connection to this computer** to ignore all future certificate warnings, you will not be able to connect to Data Protection for Microsoft Hyper-V.

- b) In the **General** tab of the **Certificate** window, click **Install Certificate**.

- c) In the welcome page of the **Certificate Import Wizard** window, select a store location (**Current User** or **Local Machine**) and click **Next**.

- d) In the **Certificate Store** page, click **Place all certificates in the following store** and click **Browse**.

- e) In the **Select Certificate Store** window, select **Trusted Root Certification Authorities** and click **OK**.
- f) Click **Next** in the **Certificate Store** page.
- g) Review the selections in the **Completing the Certificate Import Wizard** page and click **Finish**.
- h) In the **Security Warning** window, click **Yes** to install the certificate.
- i) Click **OK** in the confirmation window.

If you reject the certificate, you will not be able to connect to Data Protection for Microsoft Hyper-V.

What to do next

You can use the Data Protection for Microsoft Hyper-V Management Console to manage your virtual machine backups and monitor their status.

After a period of inactivity, your connection to Data Protection for Microsoft Hyper-V Management Console can time out. If a timeout occurs, you are prompted to enter your credentials in the **Reconnect - Session Expired** window.

Tip for running in a custom console: You can add the Data Protection for Microsoft Hyper-V Management Console to a custom Microsoft management console so that it can be run in a single console along with, for example, the Hyper-V manager and Cluster Failover manager.

1. Start an empty Microsoft Management Console by issuing the **mmc** command from a command prompt.
2. Click **File > Add/Remove Snap-in**.
3. Select **Data Protection for Microsoft Hyper-V** and click **Add**.
4. Select and add any other snap-ins, for example, select and add **Hyper-V Manager** and **Failover Cluster Manager**.
5. To name and store the .msc file, click **File > Save as**.
6. To start the custom console, run the .msc file that you saved.

Navigating the Data Protection for Microsoft Hyper-V Management Console

Use the Data Protection for Microsoft Hyper-V Management Console for daily management of backup operations. You can monitor virtual machine backup operations, run backup and restore operations, and update the configuration.

The Data Protection for Microsoft Hyper-V Management Console contains three main work areas: the navigation pane, the results pane, and the actions pane. Information about these work areas is provided.

Navigation pane

The navigation pane on the left, labeled **Data Protection for Hyper-V**, contains a tree view that shows the clusters or hosts in the Hyper-V environment. In the cluster view, children nodes of the cluster node represent each host in the cluster.

When you select a host or cluster in the navigation pane, the backup status of virtual machines (VMs) in the selected host or cluster and the history of schedule runs are displayed in the **Virtual Machines** and **Schedule History** views in the results pane. The list of available actions for the selected cluster, host, or VM is also displayed in the actions pane on the right of the management console.

Results pane

The results pane in the middle of the Data Protection for Microsoft Hyper-V Management Console shows detailed information about the virtual machine backups and the backup schedule history for a selected cluster or host.

The workspace contains two views, labeled as **Virtual Machines** and **Schedule History**. Click the corresponding tab in the results pane to display each view.

Virtual Machines view

The **Virtual Machines** view in the results pane shows the data protection status of each virtual machine (VM) in a cluster or host and the backup history for individual VMs.

You can enter all or part of a VM name in the **Filter** field to display only VMs with names that contain the text string. You can also click **Refresh** to refresh the contents in the tables.

VM table

At the top of the Virtual Machines view, you see the table of VMs in a host or cluster, and details about the last backup operation for each VM. The following data is shown in the table.

Table 8. Descriptions of columns in the VM table	
Column	Description
Name	The name of the VM.
Host	<p>When a cluster is selected in the navigation pane, the name of the current active host for the VM. If the VM status is Deleted, the host name is still shown.</p> <p>However, if the environment was upgraded from Data Protection for Microsoft Hyper-V Version 8.1.2 or earlier, the field will be empty until VM backups are run with version 8.1.27.</p>
State	<p>The state of the selected VM. A VM can have one of the following states:</p> <p>Off The VM is powered off.</p> <p>Running The VM is powered on and running.</p> <p>Paused The current state of the VM is preserved, but the memory that is allocated for the VM is still in use while the processor resources are freed.</p> <p>Saved The current state of the VM is preserved, but the memory that is allocated for the VM and the processor resources is freed.</p> <p>If no entry is shown, the VM was deleted from your environment, but the backup still exists on the IBM Storage Protect server.</p>

Table 8. Descriptions of columns in the VM table (continued)

Column	Description
Status	<p>The data protection status of the VM. A VM can have one of the following protection status:</p> <p>At Risk The most recent backup operation did not occur within the time limit that is specified by the at-risk policy.</p> <p>No Backup The VM is configured for backup operations, but no backup has been run.</p> <p>Normal A backup operation occurred within the time limit that is specified by the at-risk policy.</p> <p>Ignored The at-risk policy is set to suppress at-risk warnings for the VM.</p> <p>Deleted The VM was deleted from the Hyper-V environment, but its backup is available to be restored.</p> <p>Instant Access The VM is running in instant access mode. An instant access VM is a temporary copy of a VM backup that is used for operations such as the verification of backup data.</p>
Last Backup	The date of the last successful backup operation.
Data Transmitted	The amount of data that was sent to the IBM Storage Protect server during the backup operation.
Duration	The length of time it took to run the backup operation.
Backup Type	The type of backup operation that was run (full or incremental).
Schedule	The name of the schedule that ran during the last successful backup operation.

Backup History table

The **Backup History** table shows the details of previously scheduled or ad hoc backup tasks of a single virtual machine (VM) that you selected in the VM table. If you selected multiple VMs, no data is displayed 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 **SET EVENTRETENTION** command on the IBM Storage Protect server.

The following data is shown in the table.

Table 9. Descriptions of columns in the **Backup History** table

Column	Description
Last Run Time	The actual start date and time of the last backup run.
Status	<p>The status of the backup operation.</p> <p>Succeeded The backup operation was completed successfully.</p> <p>Failed The backup operation encountered an error and was not completed.</p> <p>In Progress A backup operation is in progress.</p>
Duration	The duration of the backup operation.
Error Code	If a backup operation failed, an error code is shown. If the backup operation completed successfully, a zero (0) is displayed.
Data Transmitted	The amount of data that was sent to the IBM Storage Protect server during the backup operation.
Backup Type	<p>The type of backup operation that was run for the VM:</p> <p>Incremental Backs up the blocks that changed since the last backup (full or incremental).</p> <p>Full Backs up a snapshot of an entire VM.</p>
Backup Host	The host that contains the data mover for the VM when it was backed up. For clusters, this data mover host can change due to failover clustering.

Tasks table

The **Tasks** table shows a list of recent tasks that started since the Data Protection for Microsoft Hyper-V Management Console was started.

For more information, see [“Tasks table” on page 68](#).

You can also show detailed information about an active task. For more information, see [“Task Details” on page 68](#).

Related tasks

[“Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console” on page 130](#)

For ease of use, you can run instant access operations from the Data Protection for Microsoft Hyper-V Management Console.

[“Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console” on page 131](#)

When you no longer need an instant access virtual machine (VM), you can use the Data Protection for Microsoft Hyper-V Management Console to remove the VM.

Schedule History view

The **Schedule History** view in the results pane displays the run history for the backup schedules that are associated with a Hyper-V host or cluster.

You can click **Refresh** to refresh the contents in the tables.

Schedule history table

The Schedule history table shows the history of backup schedules for the host or cluster.

The number of listings of backup history that are shown depends on the number of days that are set by the **SET EVENTRETENTION** command on the IBM Storage Protect server.

The following data is shown in the schedule history table.

Table 10. Descriptions of columns in the schedule history table	
Column	Description
Schedule Start Time	The actual date and time that the schedule started. If a schedule was missed, the scheduled start time is shown.
Name	The name of the schedule.
Status	<p>The status of the schedule is based on all data movers that are associated with the schedule. The following states are possible:</p> <p>Succeeded The schedule ran to completion for all data movers. The details of individual VMs that were backed up or failed to be backed up are shown in the second table.</p> <p>Failed The schedule did not run to completion on at least one data mover.</p> <p>In Progress The schedule started on all data movers and has not completed.</p> <p>Missed The schedule failed to start on at least one data mover within the startup window for the schedule.</p>
VM Succeeded	The number of VMs that were successfully backed up during the schedule run.
VM Failures	The number of VMs that failed to be backed up during the schedule run. If the schedule was missed or failed, a dash is displayed.
Duration	The length of time that the schedule ran. The duration is measured from the start of the first schedule activity to the final schedule activity. If the schedule was missed or failed, a dash is displayed.

Schedule Detail table

When you select a schedule entry in the schedule history table, the **Schedule Detail** table shows the list of virtual machines (VMs) that were backed up for the selected schedule run.

When multiple nodes are associated with a schedule, the number of virtual machines (VMs) that are shown reflects the information from all the data mover nodes for that schedule run.

You can enter all or part of a VM name in the **Filter** field to display only VMs with names that contain the text string.

The following data is shown in the table.

Table 11. Descriptions of columns in the Schedule Detail table	
Column	Description
Name	The name of the VM.
Status	The backup status of the VM. Succeeded The VM was successfully backed up. Failed The VM failed to be backed up.
Start Time	The date and time when the VM backup operation started.
Reason	If the VM backup failed, an error code is provided. If the backup operation was successful, a zero (0) is displayed.
Duration	The duration of the backup operation.
Data Transmitted	The amount of data that was sent to the IBM Storage Protect server during the backup operation.
Backup Type	The type of backup operation that was run for the VM: Incremental Backs up the blocks that changed since the last backup (full or incremental). Full Backs up a snapshot of an entire VM.
Backup Host	The host that contains the data mover that is used to run the VM backup operation. For clusters, this data mover host can change due to failover clustering.

Tasks table

The **Tasks** table shows a list of recent tasks that started since the Data Protection for Microsoft Hyper-V Management Console was started.

For more information, see [“Tasks table” on page 68](#).

You can also show detailed information about an active task. For more information, see [“Task Details” on page 68](#).

Tasks table

The **Tasks** table shows the list of recent backup or restore tasks that began since you started the Data Protection for Microsoft Hyper-V Management Console.

The same list of tasks is displayed in the **Virtual Machines** view or the **Schedule History** view. You can monitor long-running tasks such as backup or restore operations.

You can take more actions by clicking the following buttons:

Stop

Cancels a running task.

Remove Completed

Removes all completed tasks from the table. Tasks that are running are not removed.

Show Details

Opens the **Task Details** pane to show more information about an active task. The **Show Details** button is visible only when the **Task Details** pane is closed. For more information, see [“Task Details” on page 68](#).

Hide Details

Closes the **Task Details** pane. The **Hide Details** button is visible only when the **Task Details** pane is opened.

The following data is shown in the **Tasks** table.

Table 12. Descriptions of columns in the Tasks table	
Column	Description
Host	The host on which the task is running.
Task	The type of task that is running (Backup , Restore , or Restore - Instant Access).
Status	The status of the task (Working , Succeeded , or Failed).
Start Time	The start date and time of the task.
Duration	The length of time it took for the task to run or the length of time the task has been running.
Messages	<p>If the task failed, the related error messages are shown. If the task was completed successfully, no messages are shown.</p> <p>The messages field also shows status messages for a task that is in progress.</p>

Task Details

The **Task Details** pane shows detailed information about an active task that is selected in the **Tasks** table. You can also copy the details of the task to the clipboard.

You open the **Task Details** pane by selecting a task that is running in the **Tasks** table and clicking **Show Details**.

You close the **Task Details** pane by clicking **Hide Details**.

To display detailed information about a nonzero error code, click **Error Information**.

The following table describes the data that is shown in the list boxes in the **Task Details** pane.

Table 13. Description of data in the **Task Details** pane

Box name	Description
Type of operation	Shows the type of operation such as IFINCREMENTAL backup or IFFULL backup . The names of the VMs that are being processed are listed in the list box. You can click a VM name to show the details in the Status and Messages boxes.
Status	Shows the status information about a task. The information is updated while the task is running. Objects inspected The total number of objects inspected during the operation. Objects backed up The total number of objects backed up during the operation. Objects failed The total number of objects for which the operation failed during the operation. For detailed information, see the <code>dsmerror.hostname_HV_DM.log</code> file. Total bytes The total number of bytes transferred during the operation. Aggregate network The total number of bytes transferred to the IBM Storage Protect server divided by the total elapsed time during an operation. Error code Any return codes that are generated during the operation. Click Error Information to display detailed information about the return code. For more information, see “IBM Storage Protect Error Information” on page 70 .
Messages	Shows any messages that are generated during the task. Time Shows the date and time when the message was issued. Message Shows the list of messages that are issued during the task.

Tip: You can copy the task details to the clipboard by clicking **Copy**. The task details are the same as the information that is displayed in the **Task Details** pane, along with the equivalent Windows PowerShell cmdlet for the task.

For example, the following details are copied to the clipboard during an active task:

```
KINGSTON7      Backup      Working      9/5/2018 1:55:52 PM      00.00.00:40
IFINCREMENTAL backup
```

```
VM Name
kingston19
```

```
Status
Objects inspected: 1
Objects backed up: 0
Objects failed: 0
Total bytes: 0.00 Bytes
Aggregate network: 0 KB/sec
Return code: 0
```

```
Messages
9/5/2018 1:56:11 PM Comparing...
```

```
PowerShell
Backup-DpHvVM -session $sessKINGSTON7 -vmName kingsont19 -node IFINCREMENTAL
-AppConsistentAttempt 2 -CrashConsistentAttempt 0 -diskProtrection 0
```

Restriction: If you are backing up multiple VMs and you included only certain virtual disks in the **Ad Hoc Backup** window, the task details for only one VM are displayed during the backup operation, even though the operation was completed successfully for all VMs.

For example, you selected **Include disk 1 only** in the **Data Protection** field in the **Ad Hoc Backup** window. When you back up multiple VMs (VM1, VM2, and VM3) for backup operations in the Data Protection for Microsoft Hyper-V Management Console, the **Task Details** pane will show only the details for VM1 during the backup operation.

This restriction applies only to the Data Protection for Microsoft Hyper-V Management Console. Backup operations on the data mover command line are not affected.

IBM Storage Protect Error Information

If the status of a task contains a nonzero return code in the **Task Details** pane, you can click **Error Information** to display detailed information about the return code. The detailed error information is displayed in the **IBM Storage Protect Error Information** window.

The following information is available in the **Error Information** window:

- The error code.
- The error message.
- A link to a website that includes documentation for the messages.
- The name of the data mover node.
- The host name.
- The name of the error log file.
- A section of the error log for the data mover node. Displays a maximum of 200 lines, starting with the task start time up to the task end time.

You can also download the error log file or copy the section of the error log that is shown in the window:

Download Log

Downloads the error log file. When you click **Download Log**, a window is opened to indicate whether the file was saved. If the file is saved, it shows the saved file name and prompts you to open the log in **Notepad**.

Copy

Copies the section of the error log that is shown and pastes it to the clipboard.

Close

Closes the error information window.

Tip: You can also use the following PowerShell cmdlets to return detailed information about a return code:

- **Get-DpHvRCInfo**

- **Get-DpHVLogSection**
- **Get-DpHvLog**

For more information, see [“PowerShell cmdlets for Data Protection for Microsoft Hyper-V” on page 135.](#)

Actions pane

The **Actions** pane on the right side of the Data Protection for Microsoft Hyper-V Management Console shows the list of available actions for the selected item in the navigation pane and any selected VMs in the results pane.

The actions pane contains one section for a host or cluster, and one section for a VM.

Actions that apply at the host or cluster level

Log out

Log out of Data Protection for Microsoft Hyper-V.

Connect

Log in to Data Protection for Microsoft Hyper-V.

Backup Management

Assign a backup schedule to a single Hyper-V host or cluster environment.

Configure

Open the configuration wizard to update the configuration of Data Protection for Microsoft Hyper-V.

Properties

Show the current configuration for Data Protection for Microsoft Hyper-V. To update the configuration, click **Configure**.

View > Customize

Customize what is displayed in the Data Protection for Microsoft Hyper-V Management Console.

Refresh

Refresh the contents in the Data Protection for Microsoft Hyper-V Management Console.

Help

Open the online help for Data Protection for Microsoft Hyper-V Management Console.

Actions that apply at the VM level

Backup

Back up one or VMs.

Restore

Restore a single VM with the **Restore** wizard.

File Restore

Open the file restore interface in a web browser. Available only if you enabled the file restore feature.

Set At Risk

Set the at-risk policy for one or more VMs.

Help

Open the online help for Data Protection for Microsoft Hyper-V Management Console.

Verifying the configuration of Data Protection for Microsoft Hyper-V

After you run the configuration wizard, you can use the Data Protection for Microsoft Hyper-V Management Console to verify the configuration of the nodes that were created during the configuration

process. By verifying the node configuration, you can help to prevent potential issues with system operations.

About this task

When you verify the configuration of the nodes, the following types of information are displayed:

- Information about the data mover node such as the host name, operating system, and location of the error log
- If the file restore feature is enabled, information about the mount proxy nodes such as the host name, operating system, location of the error log, the state of the recovery agent, and the iSCSI status of the mount proxy nodes

Procedure

1. [Start the Data Protection for Microsoft Hyper-V Management Console.](#)
2. Select a cluster or host from the navigation pane.
3. In the **Actions** pane, click **Properties**.
4. In the **Properties** window, select the **Verify Nodes** page to display the node information.

The data that is displayed on the **General** page and **Verify Nodes** page depends on the node that you selected in the navigation pane. If you selected a cluster node, the information about all valid nodes in the cluster is displayed. If you selected a host, only the data that is related to the host is displayed.

5. Select a node that you want to verify from the **Node Information** box and click **Verify Node**.

Tip: If you select a Linux mount proxy node, the **Verify Node** button is disabled. To view Linux mount proxy information, select the Windows mount proxy node (usually the next item in the list) and click **Verify Node**. Then, select the Linux mount proxy node in the **Node Information** box again to view the Linux mount proxy information in the **Status Details** box.

6. Review the results in the **Status Details** box and resolve any issues that are discovered during the verification.

Tip: You can save the results to the clipboard by highlighting the contents in the **Status Details** box, and pressing **Ctrl+C**. You can then paste the contents into a text document and save it for reference.

7. To close the **Properties** window, click **Close**.

What to do next

After you resolve any configuration issues, you can restart the Data Protection for Microsoft Hyper-V Management Console and verify the configuration again.

Tip: You can also verify the configuration with the **Test-DpHvConfiguration** PowerShell cmdlet. For more information, see [“Data Protection for Microsoft Hyper-V cmdlet examples”](#) on page 138.

Managing backup schedules for a host or cluster machine

You can select a schedule to specify how often and when to automatically back up virtual machines (VMs) in a Hyper-V host or cluster.

About this task

Schedules are set up by the IBM Storage Protect server administrator to automatically back up VMs.

To enable schedules to be used for Data Protection for Microsoft Hyper-V, the IBM Storage Protect server administrator must set up a list of schedules that are specifically for backing up Hyper-V VMs. The schedule definition must include the following parameters and options:

- The `-domain.vmfull="all-vm"` option must be specified in the option string. No other parameters are required for the `-domain.vmfull` option.

- The schedule must contain the ACTION=BACKUP and SUBACTION=VM parameters.

For example, the administrator defines a schedule with the following **DEFINE SCHEDULE** server command:

```
define schedule hyperv_domain_name schedule_name
description=schedule_description action=backup subaction=VM
starttime=NOW+00:10 schedstyle=Classic period=1 perunits=Weeks
durunits=minutes duration=10 options='-vmbackuptype=hypervfull
-mode=IFIncremental -domain.vmfull="all-vm"'
```

The Hyper-V administrator then associates a data mover with the schedule by using the **Backup Management** window. The -asnodename= option is automatically added to the schedule definition. For example, for a stand-alone host, the following option is added:

```
-asnodename=Hyper-V_host_HV_TGT
```

For a cluster environment, the following options is added:

```
-asnodename=clustername_hv_tgt
```

Depending on the system configuration, the node name can also contain a prefix and suffix. For more information, see [“Customizing node names”](#) on page 20.

Tip: The server administrator can also use the IBM Storage Protect Operations Center to define the Hyper-V schedule.

If some VMs need to be excluded, include the -vm parameter in the -domain.vmfull option in the option string. For example, to back up all VMs but exclude the VM named TestVm1, specify the following parameters in the option string:

```
-domain.vmfull="all-vm;-vm=TestVM1"
```

To include a single VM in scheduled backup operations, specify the following parameters in the option string. In the following example, the VM is named TestVM1:

```
-domain.vmfull="all-vm;-vm=*;vm=TestVM1"
```

You set the backup policy for a host or cluster by associating a backup schedule with the host or cluster. You can also remove the schedule association from a host or cluster.

In a cluster environment, the selected schedule applies to all hosts in the cluster. You cannot assign a different schedule to a host that is part of a cluster.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#)
2. In the navigation pane, click a stand-alone host or cluster from the navigation pane.
3. In the **Actions** pane, click **Backup Management**.

A summary of the schedules is displayed in a table. The following properties of the schedules are displayed:

Schedule Name

The name of the schedule.

Repeats

How often the schedule repeats.

Host Names

A list of hosts that correspond to the data mover nodes that are associate with the schedule.

Description

A description of the schedule.

4. Select a schedule in the **Backup Management** window and take one of the following actions.

- To associate the selected schedule with the cluster or host and refresh the window, click **Assign Schedule**.

When a schedule is assigned to a cluster or host, the -asnodename in the option string in the schedule definition is specified with the target node (*hostname_HV_TGT* or *clustername_HV_TGT*). Only the relevant schedules for this target node are shown or schedules that are not associated with any other target nodes are shown.

The target node name can also contain a prefix and suffix. For example, *prefix_hostname_HV_TGT_suffix* or *prefix_clustername_HV_TGT_suffix*.

- To remove the selected schedule association from the cluster or host, click **Unassign Schedule**.

When you remove the schedule association from the cluster or host, the asnodename option is removed from the option string in the schedule definition and the nodes related to the cluster or host are removed from the association.

5. Click **Close** to close the window.

Setting the at-risk policy for a virtual machine

Hyper-V 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 whether the VM is shown as being at-risk if a backup operation does not occur in a specified time interval.

About this task

By default, each VM uses the policy that is set for the IBM Storage Protect server. You can use the default policy, set a custom policy, or choose to ignore the policy for one or more VMs that are selected in the cluster or host view in the Data Protection for Microsoft Hyper-V Management Console.

The **Status** column in the **Virtual Machines** pane shows the data protection status of each VM in the host or cluster. The following data protection states are possible.

At Risk

The most recent backup operation did not occur within the time limit that is specified by the at-risk policy.

No Backup

The VM is configured for backup operations, but no backup has been run.

Normal

A backup operation occurred within the time limit that is specified by the at-risk policy.

Ignored

The at-risk policy is set to suppress at-risk warnings for the VM.

Deleted

The VM was deleted from the Hyper-V environment, but its backup is available to be restored.

Instant Access

The VM is running in instant access mode. An instant access VM is a temporary copy of a VM backup that is used for operations such as the verification of backup data.

You can assign an at-risk policy only to VMs that have been backed up. If a VM has never been backed up, the set at-risk action is disabled.

Procedure

To use 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. Start the Data Protection for Microsoft Hyper-V Management Console
2. In the navigation pane, click a host or cluster, and click one or more VMs in the VM table.
3. Click **Action > Set At Risk**.

4. Complete one of the following actions in the **Set At Risk** window.

Action	Step
To use the default at-risk policy	Click Default to accept the default duration of 1 day.
To suppress at-risk warnings for the VM	Click Ignore .
To set a custom at-risk policy	Click Custom and move the slider to set the time interval since the last backup. The default is 6 hours.

5. To save your setting, click **Set at Risk**.

6. To close the window, click **Close**.

Results

If the at-risk policy is set to **Default** or **Custom** for a VM, the **At Risk** status is shown for the VM if a backup operation did not occur within the time limit that is set by the policy. If the VM has never been backed up, the VM is also considered to be at-risk and the **No Backup** status 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.

Viewing the schedule history for a Hyper-V host or cluster

You can view the run history for the backup schedules that are associated with a Hyper-V host or cluster. 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 number of runs that are shown for a schedule depends on the number of days that are set by the **SET EVENTRETENTION** command on the IBM Storage Protect server.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. In the navigation pane, click a host or cluster, and click the **Schedule History** tab.

You can view the run history for all of the backup schedules that are associated with the cluster or host. You can also select a schedule to view the backup status for the VMs that are associated with that schedule in the **Schedule Detail** table.

For information, see [“Schedule History view”](#) on page 66.

Related information

[SET EVENTRETENTION \(Set the retention period for event records\)](#)

Viewing the backup status and backup history of a virtual machine

You can view the status of scheduled virtual machine (VM) backups in a host or cluster to identify the VMs that might require attention. You can also view the backup history of individual VMs.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. From the navigation pane, click a host or cluster.
3. Click the **Virtual Machines** tab.
4. In the VM table, view the status of the most recent scheduled backup operations of VMs in the host or cluster.

A VM with a status of **At Risk** indicates that a last scheduled backup was missed or completed with errors.

5. To view the backup history of a VM, select a VM from the VM table.

The backup history that is specific to that VM 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 **SET EVENTRETENTION** command on the IBM Storage Protect server.

For more information, see [“Virtual Machines view” on page 63](#).

Tip: If you are using a data mover command (**dsmc**) to access information about the VM backups, specify the following options with the **dsmc** command:

- For clusters, include the following options:

```
-optfile=hostname_HV_DM.opt  
-asnodename=clustername_HV_TGT
```

- For a stand-alone host, include the following options:

```
-optfile=hostname_HV_DM.opt  
-asnodename=hostname_HV_TGT
```

Depending on the system configuration, the node name can also contain a prefix and suffix. In this case, specify the following options with the **dsmc** command:

- For clusters, include the following options:

```
-optfile=prefix_hostname_HV_DM_suffix.opt  
-asnodename=prefix_clustername_HV_TGT_suffix
```

- For a stand-alone host, include the following options:

```
-optfile=prefix_hostname_HV_DM_suffix.opt  
-asnodename=prefix_hostname_HV_TGT_suffix
```

For example, use the following command syntax to query information about VM backups on the IBM Storage Protect server:

```
dsmc query vm vmname -optfile=hostname_HV_DM.opt -asnodename=clustername_HV_TGT
```

If you do not include the **-asnodename** and **-optfile** options in the **dsmc query vm** command, the output of the command will not match the VM backup results in the Data Protection for Microsoft Hyper-V Management Console.

What to do next

If you want to back up an at-risk VM without waiting for the schedule to run, select the VM, and click **Backup** from the **Actions** pane.

Related tasks

[“Customizing node names” on page 20](#)

You can add a prefix, suffix, or both to default node names. In this way, you can customize the node names that are generated automatically by the configuration wizard.

Related information

[SET EVENTRETENTION \(Set the retention period for event records\)](#)

Running an ad hoc backup of a virtual machine

When you start an ad hoc 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 Hyper-V host or cluster are backed up when a schedule is run. However, you can start an ad hoc backup if you notice that a backup schedule was missed or if a VM backup was completed with errors. You can also start an ad hoc backup of a VM that is excluded from scheduled backup services.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. In the navigation pane, click a host or cluster.
3. In the VM table in the **Virtual Machines** view, click a VM.
For example, click a VM whose data protection status is **At Risk**.
4. Click **Backup** in the VM section of the **Actions** pane.
5. Complete the following fields in the **Ad Hoc Backup** window:

Option	Description
Backup Type	Select the type of backup to run: Incremental Backs up the blocks that changed since the previous backup (full or incremental). The most recent incremental backup 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. Full Backs up a snapshot of an entire VM. After the full backup is completed, you do not have to make additional full backups.
Data Consistency	Available only for Hyper-V hosts or clusters on Windows Server 2016 operating systems. Select the type of snapshot and retry attempts that are used during backup operations: Always application consistent Attempts two quiesced snapshots to create application-consistent backups before failing the backup. Attempt application consistent Attempts one quiesced snapshot and as a final attempt, a non-quiesced, machine-consistent snapshot. Machine consistent only Attempts only a non-quiesced snapshot for VMs that can never complete a quiesced snapshot.
Disk Protection	Select the VM disks to include in backups. The disks are identified by the disk numbers. You can back up all disks in the VM, back up only disk 1, or back up all disks except for disk 1. Disk 1 usually contains the operating system.

6. To start the backup operation and close the window, click **Backup**.

Results

The backup operation that you started is displayed in the task list at the bottom of the **Virtual Machines** view or the **Schedule History** view.

Restoring a virtual machine

You can restore a virtual machine (VM) from a backup that is on an IBM Storage Protect server.

About this task

During the restore operation, the VM is shut down and deleted before it is restored from the VM backup that is stored on the IBM Storage Protect server. The restore operation then re-creates the VM such that its content and configuration is identical to what it was when the backup occurred. Even though the VM is shut down before it is deleted, it is a good practice to manually shut down the VM before you start the restore operation to bring any in-progress application activities to an orderly stop.

You can use the Data Protection for Microsoft Hyper-V Management Console to restore data to a new VM or replace the existing VM with the restored data.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. In the navigation pane, click a host in the cluster or host view.
3. Select a VM from the VM table in the **Virtual Machines** view.
For example, click a VM whose data protection status is **Normal**.

Tip: If you need to restore a VM that was deleted but whose backup is still available on the IBM Storage Protect server, select a VM with the **Deleted** status.

4. In the **Actions** pane, click **Restore**.
5. Complete the following pages in the **Restore** wizard as applicable. The pages that are provided depend on the options that you select in the wizard.

Wizard page	Action
Before you begin	Click Next to start the wizard.
Select restore point	<p>The highlighted dates on the calendar contain restore points. Restore points are VM backups that are available for restore operations. Some VMs have more than one restore point per day.</p> <p>Select a date and a restore point from the Available restore points list. The size of the VM is listed next to an available restore point. The VM is restored to the state in which it existed when it was backed up.</p>
Select options	<p>Create a VM or replace the existing VM with the data from the selected restore point. The following options are available:</p> <p>Create a new virtual machine Create a VM with the data from the selected restore point. This option is the default.</p> <p>Replace an existing virtual machine Replace the existing VM with the data from the selected restore point. The VM identifiers are maintained.</p> <p>Virtual machine name If you are creating a VM, the default name for the new VM is the original VM name appended with the date of the restore operation. If you do not want to use the default name, enter a VM name that is not already used by another VM in the Hyper-V host or cluster in the entry field.</p>

Wizard page	Action
	<p>If you are replacing an existing VM, the original VM name is shown. You cannot update it.</p> <p>Restore virtual machine to</p> <p>If you are creating a VM, select a host to which the VM can be restored.</p> <p>If you are replacing an existing VM, the VM is restored to the host that owns the VM. This field is not selectable.</p>
Select storage	<p>The page appears only if you are creating a VM with the data from the restore point.</p> <p>Enter the location on the host where you want to create the VM. The default location is C:\ProgramData\Microsoft\Windows\Hyper-V\Virtual Machines.</p>
Summary	Review the options that you selected in the wizard. Click Next to start the restore operation.
Results	Click Finish to close the wizard.

Results

The restore operation that you started is displayed in the task list at the bottom of the **Virtual Machines** view or **Schedule History** view.

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

What to do next

For a host that is in a cluster, if you restored a deleted VM or if you restored a VM with a new VM name, you must configure the restored VM for high availability by using Microsoft Failover Cluster Manager, System Center Virtual Machine Manager, or PowerShell cmdlets. For instructions about configuring a VM for high availability, see the Microsoft documentation.

Best practices for Data Protection for Microsoft Hyper-V

You can follow best practices to take advantage of features that can help you manage Data Protection for Microsoft Hyper-V operations.

Excluding virtual machines from scheduled backup operations

If a virtual machine (VM) is undergoing maintenance operations or if it is a test VM that does not need to be backed up regularly, you can exclude it from scheduled backup operations. Instead of updating the client options file (dsm.opt) on every data mover, consider excluding the VMs in the schedule definition on the IBM Storage Protect server.

The IBM Storage Protect server administrator can accomplish this task by adding the -vm=vmname1,vmname2 parameter to the option string on the schedule definition on a server.

For example, the administrator defined the following schedule on the IBM Storage Protect server:

```
define schedule hyperv_domain_name schedule_name
description=schedule_description action=backup subaction=VM
starttime=NOW+00:10 schedstyle=Classic period=1 perunits=Weeks
durunits=minutes duration=10 options='-vmbackuptype=fullvm
-asnodename=Hyper-V_host_hv_tgt -mode=IFIncremental
-domain.vmfull="all-vm"
```

To exclude a VM named `testvm1` from scheduled backup operations, update the `-domain.vmfull` option in the schedule definition as follows:

```
define schedule hyperv_domain_name schedule_name
description=schedule_description action=backup subaction=VM
starttime=NOW+00:10 schedstyle=Classic period=1 perunits=Weeks
durunits=minutes duration=10 options='-vmbackuptype=fullvm
-asnodename=Hyper-V_host_hv_tgt -mode=IFIncremental
-domain.vmfull="all-vm;-vm=testvm1"'
```

To exclude one or more VMs that begin with the name `testvm`, update the `-domain.vmfull` option in the schedule definition as follows:

```
define schedule hyperv_domain_name schedule_name
description=schedule_description action=backup subaction=VM
starttime=NOW+00:10 schedstyle=Classic period=1 perunits=Weeks
durunits=minutes duration=10 options='-vmbackuptype=fullvm
-asnodename=Hyper-V_host_hv_tgt -mode=IFIncremental
-domain.vmfull="all-vm;-vm=testvm*"
```

For more information about the `-domain.vmfull` option, see [“Domain.vmfull” on page 169](#).

Tip: Depending on the system configuration, the node name can also contain a prefix and suffix. For more information, see [“Customizing node names” on page 20](#).

Rebinding virtual machines to management classes

If you need to override the management class that is bound to virtual machines (VMs) by the `vmmc` option in the client options file, you can use the `include.vm` option to bind the VMs to a new management class.

In a cluster environment, you must set the `include.vm` option in the options file (`dsm.opt`) on all the hosts.

For instance, you want to back up the VMs in your test environment, but you do not want the same retention period for the test VM backups as specified by the STANDARD management class. In this case, you can rebind the test VMs to a management class that has a shorter retention period for backups.

For example, to rebind all VMs with names that begin with `testvm` with the management class named `NONPRODMC`, add the following statement to the client options file (`dsm.opt`):

```
include.vm vmtest* NONPRODMC
```

For more information and examples about the `include.vm` option, see [“Vmmc” on page 206](#).

Chapter 5. Getting started with file restore operations

To restore files from a web-based interface with minimal administrator assistance, the file restore interface is available for use. After the file restore feature is configured, the administrator sends the file restore URL to file owners or help desk personnel so that they can find and restore files.

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

When the configuration is complete, no administrator interaction is needed to access or restore files. During the configuration process, the administrator gives the file owner access to the virtual machine (VM) that contains the file owner's data. The data can be accessed with local VM credentials so that administrators can monitor file restore resources. File owner permissions do not have to be managed.

In the file restore interface, all users can view demonstration videos to learn about the IBM Storage Protect file restore interface. The *Find and Restore Files* and *Monitoring Restores* videos are displayed when users initially log on to the file restore interface. Videos are available in English only.

Related concepts

[“File restore tasks” on page 81](#)

Different types of users set up and use the file restore feature. File owners, help desk personnel, and administrators are responsible for different sets of tasks.

Related tasks

[“Logging in to restore files” on page 84](#)

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

File restore tasks

Different types of users set up and use the file restore feature. File owners, help desk personnel, and administrators are responsible for different sets of tasks.

File owner

The file owner maintains business data such as text documents, spreadsheets, and presentation files on virtual machines (VMs).

The file owner completes the following tasks to restore individual files and folders:

- [“Logging in to restore files” on page 84](#)
- [“Restoring files from a virtual machine backup” on page 84](#)

Help desk personnel

Personnel in the help-desk environment assists file owners in restoring their data.

The help desk personnel provides the specific file restore URL for file owners, or restores files on behalf of the file owners.

The help desk personnel completes the following tasks:

- Obtains the file restore URL from the file restore administrator or from the Data Protection for Microsoft Hyper-V Management Console. For more information, see Step 5 of [“Enabling the environment for file restore operations” on page 44](#).
- [“Logging in to restore files” on page 84](#)

- [“Restoring files from a virtual machine backup” on page 84](#)

File restore administrator

The administrator installs software, schedules VM backup operations to the IBM Storage Protect server, and manages user accounts and permissions in the Microsoft Hyper-V environment.

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

1. [“Enabling the environment for file restore operations” on page 44](#)
2. If you expect that file owners will run file restore operations on Linux guest VMs, complete the following tasks:
 - a. [“Installing the Linux mount proxy feature” on page 33](#)
 - b. [“Configuring the Linux mount proxy for file restore operations” on page 45](#)
3. To verify that backup operations are running as expected, wait for a scheduled backup to be completed or [run an ad hoc backup operation of a VM](#).

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

- [“Modifying options for file restore operations” on page 48](#)
- [“Configuring Data Protection for Microsoft Hyper-V log activity” on page 49](#)

If you no longer have to use file restore operations, you can remove it by following the instructions:

[“Removing the file restore feature” on page 38](#)

File restore prerequisites

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

To enable the file restore feature, IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V must be installed on a Hyper-V host system.

The file restore web service and the underlying Data Protection for Microsoft Hyper-V environment must be installed, configured, and operational, including the mount proxy data movers and ISCSI services.

Hyper-V administrators must provide file owners with a URL to connect to the file restore web interface. When you use the configuration wizard to configure Data Protection for Microsoft Hyper-V and enable the file restore feature, the URL is provided at the conclusion of the wizard. For more information, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard” on page 39](#).

Hyper-V virtual machine prerequisites

The following prerequisites apply to the Hyper-V virtual machine (VM) that contains the files to be restored:

- The VM must be running during the file restore operation.
- The Windows VM must belong to the same Windows domain as the Hyper-V host where the IBM Storage Protect backup-archive client is installed.
- When a VM is deleted from a Windows domain and restored later, the VM must rejoin the domain to ensure that the domain trust relationship is reestablished. Do not attempt a file restore from the VM until the domain trust relationship is restored.
- 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 Linux guest VMs, local user authentication is required for the VM. Authentication is not available through Windows domain, Lightweight Directory Access Protocol (LDAP), Kerberos, or other network authentication methods.

- For Linux guest VMs 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. You can specify either of the following statements:

```
ChallengeResponseAuthentication yes
```

```
#ChallengeResponseAuthentication no
```

Restart the **sshd** daemon after you modify this option.

Data mover prerequisites

A specific data mover (backup-archive client) is installed on the Hyper-V host system that "moves data" from one system to another.

The Hyper-V host system must belong to the same Windows domain as the VM 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 VM disks to be accessible as restore points to the file restore interface.

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 instructions about how to set this daemon, see the appropriate Linux documentation.

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

Microsoft Windows domain account prerequisites

The following prerequisites apply to Windows domain accounts:

- A Windows domain user with local administrator authority is required to create and access the network share. The administrator enters these credentials in the Data Protection for Microsoft Hyper-V configuration wizard to enable the environment for file restore operations.
- A file owner accesses the remote VM 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 VM and restore files into the remote VM. These credentials do not require any special permissions.
- 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 operations from tape media are not supported. The preferred method is to restore files from disk storage.

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.

Before you begin

Ensure that you obtain the URL for the file restore interface from your 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 (VM) that contains your files.
For example, `myhost.mycompany.com`.
3. Enter the user account that you use to access your files.
 - For Windows guest operating systems, use the `Windows_domain_name\user_name` format.
 - For Linux guest operating systems, use the user name that you use to log on to the Linux guest VM.
4. Enter the user account password and click **Log in**.

Related tasks

[“Restoring files from a virtual machine backup” on page 84](#)
Locate your files and restore them to a preferred location.

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 6. Protecting in-guest applications

You can use Data Protection for Microsoft Hyper-V to protect Microsoft Exchange Server and Microsoft SQL Server that run inside guest virtual machines in a Microsoft Hyper-V environment.

Protecting Microsoft Exchange Server data in Hyper-V environments

For Microsoft Exchange Server workloads that are running in a Hyper-V guest virtual machine (VM), you can take application-consistent backups of the guest VM. You can then recover a database-level or mailbox-level backup in case the original database or mailbox is damaged or lost.

The following products work together to protect Microsoft Exchange Server data in a Hyper-V environment:

- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V Version 8.1.27
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Exchange Server Version 8.1.6 or later

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

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

For the software requirements for application protection of Microsoft Exchange Server, see [technote 2017347](#).

Microsoft Active Directory support: The Hyper-V host in a cluster or non-cluster environment and the guest must use the same Active Directory instance. The host and the guest cannot use separate Active Directory instances.

Installing and configuring software for application protection of Microsoft Exchange Server

To protect a guest virtual machine (VM) that hosts Microsoft Exchange Server data, you must complete installation and configuration steps on the Hyper-V host and the guest VM. Use the step-by-step instructions to help you get your environment up and running for in-guest application protection.

Before you begin

Review the software requirements in [technote 2017347](#).

About this task

The following table lists the names that are used as examples in the tasks that follow. When you use the sample commands and options, replace the examples with the appropriate Hyper-V host and guest VM names from your environment.

Type of Name	Example
Hyper-V host or cluster name	Kingston5
Name of guest VM hosting Microsoft Exchange Server	Kingston40

Complete the following steps to install, set up, and configure Data Protection for Microsoft Hyper-V and Data Protection for Microsoft Exchange Server to protect Microsoft Exchange Server data on VM guests.

Procedure

1. [“Step 1 \(Hyper-V host\): Install and configure Data Protection for Microsoft Hyper-V” on page 88.](#)
2. [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft Exchange Server” on page 89.](#)
3. [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection” on page 90.](#)
4. [“Step 4 \(Guest VM\): Restore a database” on page 94.](#)
5. [“Optional: Configuring application protection after a virtual machine name change” on page 94](#)

Step 1 (Hyper-V host): Install and configure Data Protection for Microsoft Hyper-V

Install and configure Data Protection for Microsoft Hyper-V and ensure that you can successfully back up the guest virtual machine (VM) that hosts Microsoft Exchange Server data.

Before you begin

If you are upgrading from Data Protection for Microsoft Hyper-V 8.1.2 or earlier, rename the existing Hyper-V node name on the IBM Storage Protect server to *clustername_hv_tgt* for a cluster or *hostname_hv_tgt* for a stand-alone host. The Hyper-V node name is the node name that is specified by the *asnodename* option.

For example, rename the Hyper-V node on the server to KINGSTON_HV_TGT. For more information, see [“Renaming nodes on the IBM Storage Protect server” on page 18.](#)

Ensure that communication ports are open in the firewall as described in [“Required communication ports” on page 16.](#)

Procedure

Complete the following tasks on the Hyper-V host or cluster:

1. Install Data Protection for Microsoft Hyper-V.
For instructions, see [“Running a typical installation of Data Protection for Microsoft Hyper-V” on page 24.](#)
2. Configure Data Protection for Microsoft Hyper-V by completing the configuration wizard.
For instructions, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard” on page 39.](#)
Remember: Write down the target node name as shown on the **Cluster and Hosts** wizard page or by clicking **Actions > Properties** in the Data Protection for Microsoft Hyper-V Management Console. The target node name ends with *_HV_TGT*. The target node name is required when you run the configuration wizard in Data Protection for Microsoft Exchange Server.
3. Use the Data Protection for Microsoft Hyper-V Management Console to back up the VM that is hosting Microsoft Exchange Server.
For instructions, see [“Running an ad hoc backup of a virtual machine” on page 77.](#)
4. 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 with the **dsmc backup vm** command.
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.

What to do next

When the VM is successfully backed up, proceed to [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft Exchange Server” on page 89.](#)

Related information

[Vmpreferdagpassive](#)

Step 2 (Guest VM): Install and configure Data Protection for Microsoft Exchange Server

To ensure that you can back up databases with Data Protection for Microsoft Exchange Server, complete the steps to install and configure Data Protection for Microsoft Exchange Server, and back up a Microsoft Exchange Server database.

Before you begin

Ensure that you completed the procedure in [“Step 1 \(Hyper-V host\): Install and configure Data Protection for Microsoft Hyper-V”](#) on page 88.

Ensure that Microsoft Exchange Server databases and mailboxes are hosted on Hyper-V virtual disks.

Ensure that no Microsoft Exchange Server database is hosted on physical hard disks, independent disks, or on disks that are attached directly to the guest through in-guest iSCSI.

Procedure

Complete the following steps on the guest virtual machine (VM) that is hosting Microsoft Exchange Server data:

1. Install Data Protection for Microsoft Exchange Server if it is not already installed.

For installation instructions, see the [product documentation for IBM Storage Protect for Mail](#).

Tip: Do not run the Data Protection for Microsoft Exchange Server configuration wizard as described in the installation instructions. You will run the wizard in Step “3” on page 89. If you previously installed Data Protection for Microsoft Exchange Server and ran the wizard, you must run it again in Step “3” on page 89.

2. Install the data mover feature from the Data Protection for Microsoft Hyper-V installation package.

In the installation wizard, select the advanced installation option, and then click **Install the data mover feature or mount proxy** to install the application protection support.

For more information, see [“Installing only the Data Protection for Microsoft Hyper-V data mover”](#) on page 28.

3. Open the Data Protection for Microsoft Exchange Server Management Console by clicking **Start > DP for Exchange Management Console**.

The configuration wizard opens automatically.

If the configuration wizard does not start automatically, go to the tree view in the Management Console and click **IBM Storage Protect > Dashboard > Manage > Configuration > Wizards**. Double-click **IBM Storage Protect Configuration Wizard**.

4. On the **IBM Storage Protect Node Names** page of the configuration wizard, enter the VSS requestor, Data Protection for Microsoft Exchange Server, and Hyper-V target node names in the respective fields. Ensure that the **Do not configure DP Exchange VSS Support** check box is cleared.

For example, the following table lists node names that are used in the configuration instructions.

Field name	Node name examples
VSS Requestor	KINGSTON40_VSS
Data Protection for Exchange	KINGSTON40_EXC
Hyper-V Target Node	KINGSTON5_HV_TGT

5. On the **IBM Storage Protect Server Settings** page of the configuration wizard, complete one of the following steps:

- To configure the IBM Storage Protect server by using the wizard, select **Review** or **Edit** and update the macro as needed.
- To manually configure the server, complete the following steps:

- a. On the last wizard page, click the link that opens the macro file.
- b. Update the macro file and run it, or issue the appropriate commands from the macro, adjusting the commands as required for your environment.

For example, assume that a policy domain named `fc_m_pdexc` is set up for your use. From the `C:\Program Files\Tivoli\TSM\baclient` folder, run the **dsmadmc** command and issue the following commands:

```
register node KINGSTON40_VSS T_3_m_p_P_w userid=KINGSTON40_VSS
update node KINGSTON40_VSS T_3_m_p_P_w backdelete=yes forcep=yes
```

```
register node KINGSTON40_EXC T_3_m_p_P_w domain=fc_m_pdexc
userid=KINGSTON40_EXC
update node KINGSTON40_EXC T_3_m_p_P_w backdelete=yes domain=fc_m_pdexc
forcep=yes
```

```
grant proxynode target=KINGSTON40_EXC agent=KINGSTON40_VSS
```

The `forcep=yes` option forces the password to be reset upon first access.

In some cases, you might see the following error message when you run the **dsmadmc** command:

```
ANS1592E Failed to initialize SSL protocol
```

If this message is displayed, ensure that the `sessionsecurity` option is set to **transitional** on the IBM Storage Protect server administrator account that you are using.

For example, issue the following command from a remote computer that can access the IBM Storage Protect server:

```
update admin myAdmin sessionsecurity=transitional
```

6. Complete the configuration wizard.
7. Back up a database from the Data Protection for Microsoft Exchange Server Management Console:
 - a) In the **Actions** pane, click **Backup Destination > TSM**.
 - b) In the **Actions** pane, click **Full Backup**.
8. Optional: To ensure that consistent location information exists for the mailbox history and the mailboxes in the database backup, manually update mailbox history information.

For instructions, see [“Updating mailbox information in Microsoft Exchange Server backups”](#) on page 97.

What to do next

When the VSS backup is completed successfully, continue to [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection”](#) on page 90.

To avoid possible restore issues for the data on the guest VM, review the information in [“Shadow copy considerations for restoring an application protection backup from the data mover”](#) on page 180.

Step 3 (Hyper-V host): Configure Data Protection for Microsoft Hyper-V for application protection

Configure Data Protection for Microsoft Hyper-V to protect the guest virtual machine (VM) that is hosting Microsoft Exchange Server data. Back up the VM and verify that the backup operation was completed successfully.

Before you begin

Ensure that you completed the procedure in [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft Exchange Server”](#) on page 89.

Ensure that the virtual hard disks (VHDXs) that host the Microsoft Exchange Server database are not being excluded from the VM backup operation. For instructions, see [“Verifying that Microsoft Exchange Server volumes are not excluded in virtual machine backups”](#) on page 98.

Optional: The Integration Services or the Guest Service Interface is automatically enabled for the guest VM during a backup operation. You do not have to enable it manually. However, if you want to review the current status or enable the **Guest Services** service explicitly, use one of the following methods on the Hyper-V host or cluster:

- From Hyper-V Manager:
 1. Right-click the VM and click **Settings > Integration Services**.
 2. In the **Integration Services** window, ensure that the **Guest Services** check box is selected.
- As Administrator, issue the following commands at a PowerShell command prompt:

```
Get-VMIntegrationService -VMName Kingston40 `
Enable-VMIntegrationService -VMName Kingston40 -Name "Guest Service Interface"
```

About this task

Data Protection for Microsoft Hyper-V provides application consistency when you back up VMs that are hosting Microsoft Exchange Servers. With these backups, you can recover the VM in a consistent state.

To restore only selected databases or mailboxes from this type of backup without having to recover the entire VM, information about the state of the Microsoft Exchange Server must be preserved at the time of the VM snapshot and backup. This information is collected as part of the Microsoft Volume Shadow Copy Services (VSS) interaction that occurs during a VM snapshot.

Shadow copy considerations for restoring an application protection backup:

Typically, the shadow storage is part of a volume. However, sometimes the shadow storage space is configured on a different volume either by default or as a result of manual configuration. In this case, the database restore operation might fail because the shadow copy that was created during the VM backup operation is not available at restore time.

Before you start a backup operation, review the information in [“The shadow copy must be available on the restored volume during an application protection restore”](#) on page 181.

For Data Protection for Microsoft Hyper-V to collect the Microsoft VSS metadata for Microsoft Exchange Server, you must configure Data Protection for Microsoft Hyper-V to obtain this information from the VM during backup operations.

In this procedure, Kingston40 is an example of the Hyper-V host as shown in Microsoft Hyper-V Manager. Kingston5 is an example of the name of the guest VM that is hosting the Microsoft Exchange Server. Kingston40_vss is the name of the client node that is created during the installation of Data Protection for Microsoft Exchange Server. Replace them with the appropriate values from your environment.

Procedure

Complete the following steps on the Hyper-V host or cluster unless instructed otherwise:

1. From the `baclient` folder, set the Windows credentials for the guest VM that is hosting Microsoft Exchange Server:
 - To set credentials for a specific VM, issue the following command at the `baclient` folder from the command prompt:

```
dsmc set password -type=vmguest guest_VM_name "guest_admin_ID" guest_admin_pw
-optfile=dsm.hostname_HV_DM.opt
```

- To set credentials for all VMs that do not have specific credentials set:

```
dsmc set password -type=vmguest allvm "guest_admin_ID" guest_admin_pw
-optfile=dsm.hostname_HV_DM.opt
```

where:

vmname

The name of the guest VM that hosts Microsoft Exchange Server. The name is the VM name that is displayed in Hyper-V Manager.

guest_admin_ID

The administrator ID for the guest VM. The *guest_admin_ID* can be a Windows domain account or a local account. For example:

- For a domain account, use the *domain\username* format.
- For a local account, use the *username* format.

guest_admin_pw

The password for the administrator ID for the guest VM.

hyperv_hostname

The name of the Hyper-V host or cluster.

For example:

```
dsmc set password -type=vmguest Kingston40 "world\alan" "@lanPwd!"
-optfile=dsm.KINGSTON5_HV_DM.opt
```

The accounts that are used in the **set password** command must be valid on both the Hyper-V host or cluster and the guest VM that hosts Exchange Server data.

The **dsmc set password** command stores the guest VM credentials, which are encrypted on the system that hosts the data mover. The following minimum permissions are required for the guest VM administrator ID and password:

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

2. From the `baclient` folder, configure the data mover options file for application protection:

- a) Open the data mover options file (`dsm.hostname_HV_DM.opt`) for editing.

For example, issue the following command:

```
notepad dsm.KINGSTON5_HV_DM.opt
```

- b) Add the `include.vmtsmvss guest_vm_name` option. The *guest_vm_name* parameter can contain wildcard characters.

For example:

```
include.vmtsmvss Kingston40
```

Or

```
include.vmtsmvss Kingston*
```

3. Complete the following steps on the guest VM to enable database backups to appear in Data Protection for Microsoft Exchange Server.

- a) Generate the credentials file in the guest by running the following command at a PowerShell command prompt and enter the domain user name (*domain name\user name*) and password when prompted.

```
Get-Credential | Export-Clixml -Path `
"C:\program files\Tivoli\TSM\baclient\dsmcreds.xml"
```

The domain user must have Exchange restore permission.

- b) Verify the credentials by running the following commands in the guest VM at an Exchange Management Shell:

```
$cred = Import-CliXml -Path `
"C:\program files\Tivoli\TSM\baclient\dsmcreds.xml"
$Session = New-PSSession -Credential $cred -ConfigurationName `
Microsoft.Exchange -ConnectionUri `
http://<Exchange_server_name>/PowerShell?serializationLevel=Full `
-Authentication Kerberos
Import-PSSession -Session $Session
Get-MailboxDatabase -Server <Exchange_server_name>
```

The list of mailbox databases will be displayed correctly.

4. From the `baclient` folder on the Hyper-V host, back up the guest VM by issuing the **dsmc backup vm** command:

For example:

```
dsmc backup vm Kingston40 -optfile=dsm.KINGSTON5_HV_DM.opt
-asnode=KINGSTON5_HV_TGT
```

5. Verify the backup operation by running the **dsmc query vm** command. VM names are case-sensitive.

For example, issue the following command from the `baclient` folder:

```
dsmc query vm Kingston40 -optfile=dsm.KINGSTON5_HV_DM.opt
-asnode=KINGSTON5_HV_TGT -detail
```

The output includes text that is similar to the following example (although your version of Microsoft Exchange Server might be different):

```
Application protection type: TSM VSS
Application(s) protected: Microsoft Exchange Server 2016
```

6. From the `baclient` folder, use the **dsmc set access** command to allow the VSS node on the guest VM to access and restore the VM backup.

For example, issue the following command:

```
dsmc set access backup -type=vm Kingston40 kingston40_vss
-nodename=KINGSTON5_HV_TGT -optfile=dsm.KINGSTON5_HV_DM.opt
```

where `Kingston40` is the name of the guest VM that is shown in Microsoft Hyper-V manager and `kingston40_vss` is an IBM Storage Protect node that is created during the installation of Data Protection for Microsoft Exchange Server. This node name is in the file `C:\Program Files\Tivoli\TSM\baclient\dsm.opt` on the `Kingston40` VM.

Issue the following **dsmc query access** command to display the VM backups that the node has access to. For example:

```
dsmc query access -nodename=KINGSTON5_HV_TGT -optfile=dsm.KINGSTON5_HV_DM.opt
```

What to do next

When the backup operation is completed successfully, proceed to [“Step 4 \(Guest VM\): Restore a database” on page 94](#).

Related reference

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

[“INCLUDE.VMTSMVSS” on page 179](#)

Related information

[Set Access](#)

[Set Password](#)

Step 4 (Guest VM): Restore a database

To verify that you correctly configured application protection, restore a database with Data Protection for Microsoft Exchange Server.

Before you begin

Ensure that you completed the procedure in [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection”](#) on page 90.

Ensure that the following required are running on the guest virtual machine (VM):

1. At the command prompt, issue the following command:

```
services.msc
```

2. Locate **IBM Storage Protect Recovery Agent Service** in the list of services and, if necessary, start the service.
3. Locate **Microsoft iSCSI Initiator Service** in the list of services. If necessary, change the startup type to **Automatic** and start the service.

Procedure

Complete the following steps on the guest VM:

1. Start the Data Protection for Microsoft Exchange Server Management Console.
2. Select an **Exchange Server** instance in the tree.
3. Navigate to the **Recover** tab, and click **Refresh**.
4. Select a database entry that has the **VMVSS** backup method.
5. In the **Action** pane, click **Restore**.
6. When the restore operation is complete, review the database and any related mailboxes.

What to do next

You can now manage backups and recover data if necessary. For more information, see:

- [“Managing backup operations”](#) on page 96
- [“Restoring data”](#) on page 99

If you change the name of the guest VM after you completed the configuration steps for application protection, you must reconfigure the software with the new VM name. For instructions, see [“Optional: Configuring application protection after a virtual machine name change”](#) on page 94.

Optional: Configuring application protection after a virtual machine name change

If you changed the name of the guest virtual machine (VM) after you completed the application protection configuration, you must reconfigure Data Protection for Microsoft Hyper-V with the renamed guest VM.

Before you begin

Ensure that you have installed and configured the software to protect guest VMs that host Microsoft Exchange Server.

About this task

Complete this task only if you changed the name of a guest VM that is protected by application protection.

Procedure

1. On the data mover on the Hyper-V host or cluster, issue the following command:

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

where:

new_vmguest_displayname

The name of the new guest VM that hosts Microsoft Exchange Server. The name is the VM name that is displayed in Hyper-V Manager.

guest_admin_ID

The administrator ID for the new guest VM. The *guest_admin_ID* must be a Windows domain account or a local account. For example:

- For a domain account, use the *domain\username* format.
- For a local account, use the *username* format.

guest_admin_pw

The password for the administrator ID for the new guest VM.

2. In the data mover options file (*dsm.hostname_HV_DM.opt*), update the *include.vmtsmvss* option as follows:

```
include.vmtsmvss new_vmguest_displayname
```

where *new_vmguest_displayname* is the display name of the new guest VM in Hyper-V Manager. You can use wildcard characters.

3. From the *baclient* folder on the data mover, back up the new guest VM by using the Data Protection for Microsoft Hyper-V Management Console or by issuing the following command at the command prompt:

```
dsmc backup vm new_vmguest_displayname -optfile=dsm.hostname_HV_DM.opt  
-asnode=hostname_HV_TGT
```

4. From the *baclient* folder, use the **dsmc set access** command to allow the VSS node on the new guest VM to access and restore the VM backup.
For example, issue the following command:

```
dsmc set access backup -type=VM new_vmguest_displayname vss_requestor_node  
-optfile=dsm.hostname_HV_DM.opt
```

where:

new_vmguest_displayname

The display name of the new guest VM in Hyper-V Manager.

vss_requestor_node

The name of the VSS requester that was configured on Data Protection for Microsoft Exchange Server.

hostname

The name of the Hyper-V host or cluster where Data Protection for Microsoft Hyper-V is installed.

5. Optional: Because the VSS requester node already has access to the VMs that were backed up under the old VM name, the Data Protection for Microsoft Exchange Server Management Console shows the VMVSS databases that were backed up from both the old VM and the new VM.

If you do not want to access the VM backups with the old VM name, you must delete the access to the old VM backups. Issue the following command from the data mover on the Hyper-V host:

```
dsmc del access -optfile=dsm.hostname_HV_DM.opt -asnode=hostname_HV_TGT
```

An access list is displayed. Enter the index that corresponds to the item that you want to delete from the access list.

If you want to save space on the IBM Storage Protect server, you can delete the file space that contains the backup data for the old VM by issuing the **dsmc delete filespace** command.

Important: When you delete a file space, you delete all backup versions within that file space and you can no longer restore the data. Verify that the old VM backups are obsolete before you delete them.

Managing backup operations

After you configure the environment to protect Microsoft Exchange Server data, you can schedule virtual machine (VM) backups and separately, you can update the mailbox information in Exchange Server database backups on the VM.

Scheduling virtual machine backups

To ensure that your data is protected regularly, schedule virtual machine (VM) backups.

Before you begin

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

By default, the maximum size allowed for a virtual hard disk (VHDX) in a backup operation is 2 TB. However, you can increase the maximum size to 8 TB by using the `vmmxvirtualdisks` option. For more information, see [“Vmmxvirtualdisks”](#) on page 204.

About this task

During backup processing, Data Protection for Microsoft Hyper-V 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 VM backups and are not available to restore.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. In the navigation pane, click a stand-alone host or cluster from the navigation pane.
3. In the **Actions** pane, click **Backup Management**.
4. Select a schedule in the **Backup Management** window and click **Assign Schedule**.
5. To close the window, click **Close**.
6. Verify that the source of the schedule includes the VMs that are hosting Microsoft Exchange Server.
7. Verify that one of the following services is running:
 - If you are using a scheduler that is managed by a client acceptor (CAD), ensure that the client acceptor service is running on the data mover.
 - If you are using the stand-alone scheduler, ensure that the scheduler service is running.

Related tasks

[“Managing backup schedules for a host or cluster machine”](#) on page 72

You can select a schedule to specify how often and when to automatically back up virtual machines (VMs) in a Hyper-V host or cluster.

Updating mailbox information in Microsoft Exchange Server backups

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

About this task

Unless Data Protection for Microsoft Exchange Server is installed on the VM, mailbox history information is not automatically updated in Exchange Server database backup operations. The automatic uploading of mailbox history can also be disabled by specifying the `VMBACKUPMAILBOXHISTORY No` in the `dsm.opt` file.

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 VMs 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

Complete the following steps on the guest VM that is hosting Exchange Server data:

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 the 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.

Procedure

1. From the Microsoft Management Console (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 Microsoft Hyper-V are identified with the `vmvss` backup method.

Verifying that Microsoft Exchange Server volumes are not excluded in virtual machine backups

The volumes in Hyper-V virtual hard disks (VHDXs) must contain the Microsoft Exchange Server databases that are not excluded from the Data Protection for Microsoft Hyper-V backup processing.

Before you begin

Ensure that the databases are not on any of the following types of disks:

- Physical disks
- Independent disks
- Disks that are attached directly to the guest operating system through iSCSI.

Procedure

1. Ensure that any `EXCLUDE.VMDISK` statements in the options file on the Data Protection for Microsoft Hyper-V data mover that is used to back up the virtual machine (VM) do not inadvertently exclude VHDXs that are hosting volumes that contain Microsoft Exchange Server files, file space, database, and mailboxes.

For example:

- The `kingston40.vhdx` file contains logical volume C:
 - The `kingston40.vhdx` file contains logical volumes E: and F:
 - The disk location (IDE controller number and device location) for `kingston40_1.vhdx` is IDE 1 0.
 - The disk location for `kingston40_2.vhdx` is IDE 1 1.
 - The Microsoft Exchange Server database files to be backed up are on the E: and F: drive.
2. Verify that no statements exclude `kingston40_2.vhdx` from the VM backup by ensuring that the data mover does not contain the following or similar statements:

```
EXCLUDE.VMDISK KINGSTON40 "IDE 1 1"
```

```
EXCLUDE.VMDISK * "IDE 1 1"
```

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

```
INCLUDE.VMDISK KINGSTON40 "IDE 1 1"
```

```
INCLUDE.VMDISK * "IDE 1 1"
```

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 VM disk from the command-line interface:

```
dsmc backup vm "KINGSTON40:-vmdisk=IDE 1 1" -asnode=KINGSTON5_HV_TGT
```

Related reference

[“Exclude.vmdisk” on page 172](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

[“Include.vmdisk” on page 176](#)

The **INCLUDE .VMDISK** option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

Restoring data

After you back up a virtual machine with the application protection feature enabled, you can recover a database or mailbox in case the original is lost or damaged.

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

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

Complete the following steps in Windows **Services**.

1. In the **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 (VM) backup by using the Data Protection for Microsoft Exchange Server graphical user interface.

Before you begin

Ensure that the Microsoft iSCSI Initiator Service is running before you restore any Microsoft Exchange Server database with backup method "VMVSS". If the service is not running, start it. For instructions, see ["Starting the Microsoft iSCSI Initiator Service" on page 99](#).

Procedure

1. To start a full database recovery from a VM, start the Data Protection for Microsoft Exchange Server Management Console (MMC).
2. In the navigation pane, expand the **Protect and Recover** node and select a Microsoft Exchange Server server.
3. On the **Recover** tab, select **Database Restore**. All backups, including all database backups from a VM backup, are listed.
4. Select a full database backup to restore.
5. 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 (VM) on the 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 VMs 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.

For instructions, see:

- [“Step 1 \(Hyper-V host\): Install and configure Data Protection for Microsoft Hyper-V” on page 88](#)
- [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection” on page 90](#)

2. On the Hyper-V host, back up vm1 and vm2 by issuing the **dsmc backup vm** command on the data mover command-line interface.
3. On vm2, install Data Protection for Microsoft Exchange Server and configure the software for Exchange Server data protection in a Hyper-V environment.

For instructions, see [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft Exchange Server” on page 89](#).

4. On the Hyper-V host, 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.

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.

For instructions on restoring mailbox data, restoring relocated and deleted mailboxes, and restoring mailbox messages interactively with Mailbox Restore Browser, see the [IBM Storage Protect for Mail product documentation](#).

Restoring data by using the command-line interface

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

Before you begin

Ensure that the Microsoft iSCSI Initiator Service is running before you restore any Microsoft Exchange Server database with backup type "VMVSS". If the service is not running, start it. For instructions, see [“Starting the Microsoft iSCSI Initiator Service” on page 99](#).

Procedure

1. 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 6.0 (C) Copyright
IBM Corporation 1997, 2018. All rights reserved.

...

Querying IBM Storage Protect server for a list of
data backups, please wait...

Connecting to IBM Storage Protect Server as node 'KINGSTON40_EXC'...
Connecting to Local DSM Agent 'exc'...
Using backup node 'KINGSTON40_EXC'...

    Exchange Server      : exc
    Database             : exc_db10

    Backup Date   Size   S   Type   Loc   Object Name
    -----
07/15/2018 19:17:26   5.40 B   A   full   Srv   20180715191726 (VMVSS)

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

2. To restore the database without applying transaction logs, issue the following command:

```
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 0.6 (C) Copyright
IBM Corporation 1997, 2018. All rights reserved.

Connecting to IBM Storage Protect Server as node 'KINGSTON40_EXC'...

Connecting to Local DSM Agent 'exc'...
Using backup node 'KINGSTON40_EXC'...

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=20180715191726 /BACKUPDEstination=TSM  
/BACKUPMethod=VMVSS
```

Restoring data by using Windows PowerShell cmdlets

If you prefer, you can use Windows PowerShell cmdlets with Data Protection for Microsoft Exchange Server to start a full Microsoft Exchange Server database recovery from a virtual machine.

Before you begin

Ensure that the Microsoft iSCSI Initiator Service is running before you restore any Microsoft Exchange Server database with backup type "VMVSS". If the service is not running, start it. For instructions, see [“Starting the Microsoft iSCSI Initiator Service” on page 99](#).

Procedure

Complete the following steps on the guest VM:

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 (VM) files. However, if the underlying file structure interests you, Data Protection for Microsoft Hyper-V backups are stored under the node name of the Hyper-V target node (for example, KINGSTON5_HV_TGT) on the IBM Storage Protect server.

The following example shows the file space information for the VM that is called Kingston40.


```
Protect: ORION>q file KINGSTON5_HV_TGT f=d

Node Name: KINGSTON5_HV_TGT
Filespace Name: \VMFULL-kingston40
Hexadecimal Filespace Name:
FSID: 61
Collocation Group Name:
Platform: TDP Hyper-V
Filespace Type: API:TSMVM
Is Filespace Unicode?: No
Capacity: 0 KB
Pct Util: 0.0
Last Backup Start Date/Time: 03/13/2018 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:
```

Protecting Microsoft SQL Server data in Hyper-V environments

For Microsoft SQL Server workloads that are running in a Hyper-V guest virtual machine (VM), you can take application-consistent backups of the guest VM. You can then recover a database-level backup in case the original database is damaged or lost.

The following products work together to protect Microsoft SQL Server data in a Hyper-V environment:

- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V version 8.1.27
- IBM Storage Protect for Virtual Environments: Data Protection for Microsoft SQL Server version 8.1.6 or later

For permissions that are required to back up and restore application data for Microsoft SQL Server, see [technote 1647995](#).

For the software requirements for application protection of Microsoft SQL Server, see [technote 2017347](#).

Installing and configuring software for application protection of Microsoft SQL Server

To protect a guest virtual machine (VM) that hosts Microsoft SQL Server data, you must complete installation and configuration steps on the Hyper-V host and the guest VM. Use the step-by-step instructions to help you get your environment up and running for in-guest application protection.

Before you begin

Review the software requirements in [technote 2017347](#).

About this task

The following table lists the names that are used as examples in the tasks that follow. When you use the sample commands and options, replace the examples with the appropriate Hyper-V host and guest VM names from your environment.

Type of Name	Example
Hyper-V host or cluster name	Kingston5
Name of guest VM hosting Microsoft SQL Server	Kingston40

Complete the following steps to install, set up, and configure Data Protection for Microsoft Hyper-V and Data Protection for Microsoft SQL Server to protect Microsoft SQL Server data on VM guests.

Procedure

1. [“Step 1 \(Hyper-V host\): Install and configure Data Protection for Microsoft Hyper-V” on page 104.](#)
2. [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft SQL Server” on page 105.](#)
3. [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection” on page 107.](#)
4. [“Step 4 \(Guest VM\): Restore a database” on page 110.](#)
5. [“Optional: Configuring application protection after a virtual machine name change” on page 110](#)

Step 1 (Hyper-V host): Install and configure Data Protection for Microsoft Hyper-V

Install and configure Data Protection for Microsoft Hyper-V and ensure that you can successfully back up the guest virtual machine (VM) that hosts Microsoft SQL Server data.

Before you begin

If you are upgrading from Data Protection for Microsoft Hyper-V Version 8.1.2 or earlier, rename the existing Hyper-V node name on the IBM Storage Protect server to *clustername_hv_tgt* for a cluster or *hostname_hv_tgt* for a stand-alone host. The Hyper-V node name is the node name that is specified by the *asnodename* option.

For example, rename the Hyper-V node on the server to KINGSTON_HV_TGT. For more information, see [“Renaming nodes on the IBM Storage Protect server” on page 18.](#)

Ensure that communication ports are open in the firewall as described in [“Required communication ports” on page 16.](#)

Procedure

Complete the following tasks on the Hyper-V host or cluster:

1. Install Data Protection for Microsoft Hyper-V.

For instructions, see [“Running a typical installation of Data Protection for Microsoft Hyper-V” on page 24.](#)

2. Configure Data Protection for Microsoft Hyper-V by completing the configuration wizard.

For instructions, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard” on page 39.](#)

Remember: Write down the target node name as shown on the **Cluster and Hosts** wizard page or by clicking **Actions > Properties** in the Data Protection for Microsoft Hyper-V Management Console. The target node name ends with *_HV_TGT*. The target node name is required when you run the configuration wizard in Data Protection for Microsoft SQL Server.

3. Use the Data Protection for Microsoft Hyper-V Management Console to back up the VM that is hosting Microsoft SQL Server.

For instructions, see [“Running an ad hoc backup of a virtual machine” on page 77.](#)

What to do next

When the VM is successfully backed up, proceed to [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft SQL Server” on page 105.](#)

Step 2 (Guest VM): Install and configure Data Protection for Microsoft SQL Server

To ensure that you can back up databases with Data Protection for Microsoft SQL Server, complete the steps to install and configure Data Protection for Microsoft SQL Server, and back up a Microsoft SQL Server database.

Before you begin

Ensure that you completed the procedure in [“Step 1 \(Hyper-V host\): Install and configure Data Protection for Microsoft Hyper-V” on page 104.](#)

Ensure that Microsoft SQL Server databases are hosted on Hyper-V virtual disks.

Ensure that no Microsoft SQL Server database is hosted on physical disks, independent disks, or disks that are attached directly to the guest through in-guest iSCSI.

Ensure that Microsoft SQL Server databases are on a single server and are not participating in any type of clustering, such as: Failover clusters, AlwaysOn Availability Groups, or AlwaysOn Failover Cluster instances.

Procedure

Complete the following steps on the guest virtual machine (VM) that is hosting Microsoft SQL Server data:

1. Install Data Protection for Microsoft SQL Server if it is not already installed.

For installation instructions, see the [product documentation for IBM Storage Protect for Databases.](#)

Tip: Do not run the Data Protection for Microsoft SQL Server configuration wizard as described in the installation instructions. You will run the wizard in Step [“3” on page 105.](#) If you previously installed Data Protection for Microsoft SQL Server and ran the wizard, you must run it again in Step [“3” on page 105.](#)

2. Install the data mover feature from the Data Protection for Microsoft Hyper-V installation package.

In the installation wizard, select the advanced installation option, and then click **Install the data mover feature or mount proxy** to install the application protection support.

For more information, see [“Installing only the Data Protection for Microsoft Hyper-V data mover” on page 28.](#)

3. Open the Data Protection for Microsoft SQL Server Management Console by clicking **Start > DP for SQL Management Console.**

The configuration wizard opens automatically.

If the configuration wizard does not start automatically, go to the tree view in the Management Console and click **IBM Storage Protect > Dashboard > Manage > Configuration > Wizards.** Double-click **IBM Storage Protect Configuration Wizard.**

4. On the **IBM Storage Protect Node Names** page of the configuration wizard, enter the VSS requestor, Data Protection for Microsoft SQL Server, and Hyper-V target node names in the respective fields. Ensure that the **Do not configure DP SQL VSS Support** check box is cleared.

For example, the following table lists node names that are used in the configuration instructions.

Field name	Node name examples
VSS Requestor	KINGSTON40_VSS
Data Protection for SQL	KINGSTON40_SQL

Field name	Node name examples
Hyper-V Target Node	KINGSTON5_HV_TGT

5. On the **IBM Storage Protect Server Settings** page of the configuration wizard, complete one of the following steps:

- To configure the IBM Storage Protect server by using the wizard, select **Review** or **Edit** and update the macro as needed.
- To manually configure the server, complete the following steps:
 - On the last wizard page, click the link that opens the macro file.
 - Update the macro file and run it, or issue the appropriate commands from the macro, adjusting the commands as required for your environment.

For example, assume that a policy domain named `fc_m_pdsq1` is set up for your use. From the `C:\Program Files\Tivoli\TSM\baclient` folder, run the **dsmdmc** command and issue the following commands:

```
register node KINGSTON40_VSS T_3_m_p_P_w userid=KINGSTON40_VSS
update node KINGSTON40_VSS T_3_m_p_P_w backdelete=yes forcep=yes
```

```
register node KINGSTON40_SQL T_3_m_p_P_w domain=fc_m_pdsq1
userid=KINGSTON40_SQL
update node KINGSTON40_SQL T_3_m_p_P_w backdelete=yes domain=fc_m_pdsq1
forcep=yes
```

```
grant proxynode target=KINGSTON40_SQL agent=KINGSTON40_VSS
```

The `forcep=yes` option forces the password to be reset upon first access.

In some cases, you might see the following error message when you run the **dsmdmc** command:

```
ANS1592E Failed to initialize SSL protocol
```

If this message is displayed, ensure that the `sessionsecurity` option is set to **transitional** on the IBM Storage Protect server administrator account that you are using.

For example, issue the following command from a remote computer that can access the IBM Storage Protect server:

```
update admin myAdmin sessionsecurity=transitional
```

- Complete the configuration wizard.
- Verify that policies are set to keep sufficient versions of Microsoft SQL Server logs and VM backups.

For instructions, see [“Managing versions of backups” on page 113](#).

- Back up a database from the Data Protection for Microsoft SQL Server Management Console:
 - In the **Actions** pane, click **Backup Method > VSS**.
 - In the **Actions** pane, click **Backup Destination > TSM**.
 - In the **Actions** pane, click **Full Backup**.

What to do next

When the VSS backup is completed successfully, continue to [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection” on page 107](#).

To avoid possible restore issues for the data on the guest VM, review the information in [“Shadow copy considerations for restoring an application protection backup from the data mover” on page 180](#).

Step 3 (Hyper-V host): Configure Data Protection for Microsoft Hyper-V for application protection

Configure Data Protection for Microsoft Hyper-V to protect the guest virtual machine (VM) that is hosting Microsoft SQL Server data. Back up the VM and verify that the backup operation was completed successfully.

Before you begin

Ensure that you completed the procedure in [“Step 2 \(Guest VM\): Install and configure Data Protection for Microsoft SQL Server”](#) on page 105.

Ensure that the virtual hard disks (VHDXs) that host the Microsoft SQL Server database are not being excluded from the VM backup operation. For instructions, see [“Verifying that Microsoft SQL Server volumes are not excluded in virtual machine backups”](#) on page 114.

Optional: The Integration Services or the Guest Service Interface is automatically enabled for the guest VM during a backup operation. You do not have to enable it manually. However, if you want to review the current status or enable the **Guest Services** service explicitly, use one of the following methods on the Hyper-V host or cluster:

- From Hyper-V Manager:
 1. Right-click the VM and click **Settings > Integration Services**.
 2. In the **Integration Services** window, ensure that the **Guest Services** check box is selected.
- As Administrator, issue the following commands at a PowerShell command prompt:

```
Get-VMIntegrationService -VMName Kingston40 `
Enable-VMIntegrationService -VMName Kingston40 -Name "Guest Service Interface"
```

About this task

Data Protection for Microsoft Hyper-V provides application consistency when you back up VMs that are hosting Microsoft SQL Servers. With these backups, you can recover the VM in a consistent state.

To restore only selected databases from this type of backup without having to recover the entire VM, information about the state of the Microsoft SQL Server must be preserved at the time of the VM snapshot and backup. This information is collected as part of the Microsoft Volume Shadow Copy Services (VSS) interaction that occurs during a VM snapshot.

Shadow copy considerations for restoring an application protection backup: Typically, the shadow storage is part of a volume. However, sometimes the shadow storage space is configured on a different volume either by default or as a result of manual configuration. In this case, the database restore operation might fail because the shadow copy that was created during the VM backup operation is not available at restore time.

Before you start a backup operation, review the information in [“The shadow copy must be available on the restored volume during an application protection restore”](#) on page 181.

For Data Protection for Microsoft Hyper-V to collect the Microsoft VSS metadata for Microsoft SQL Server, you must configure Data Protection for Microsoft Hyper-V to obtain this information from the VM during backup operations.

In this procedure, Kingston40 is an example of the Hyper-V host as shown in Microsoft Hyper-V Manager. Kingston5 is an example of the name of the guest VM that is hosting the Microsoft SQL Server. Kingston40_vss is the name of the client node that is created during the installation of Data Protection for Microsoft SQL Server. Replace the examples with the appropriate values from your environment.

Procedure

Complete the following steps on the Hyper-V host or cluster:

1. From the `baclient` folder, set the Windows credentials for the guest VM that is hosting Microsoft SQL Server:

- To set credentials for a specific VM, issue the following command at the `baclient` folder from the command prompt:

```
dsmc set password -type=vmguest guest_VM_name "guest_admin_ID" guest_admin_pw  
-optfile=dsm.hostname_HV_DM.opt
```

- To set credentials for all VMs that do not have specific credentials set:

```
dsmc set password -type=vmguest allvm "guest_admin_ID" guest_admin_pw  
-optfile=dsm.hostname_HV_DM.opt
```

where:

vmname

The name of the guest VM that hosts Microsoft SQL Server. The name is the VM name that is displayed in Hyper-V Manager.

guest_admin_ID

The administrator ID for the guest VM. The *guest_admin_ID* can be a Windows domain account or a local account. For example:

- For a domain account, use the *domain\username* format.
- For a local account, use the *username* format.

guest_admin_pw

The password for the administrator ID for the guest VM.

hostname

The name of the Hyper-V host or cluster.

For example:

```
dsmc set password -type=vmguest Kingston40 "world\alan" "@lanPwd!"  
-optfile=dsm.KINGSTON5_HV_DM.opt
```

The accounts that are used in the **set password** command must be valid on both the Hyper-V host or cluster and the guest VM that hosts SQL Server data.

Restriction: The login credentials (user name and password) for the guest VM must be the same as the credentials for the Hyper-V host.

The **dsmc set password** command stores the guest VM credentials, which are encrypted on the system that hosts the data mover. The following minimum permissions are required for the guest VM administrator ID and password:

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.

2. From the `baclient` folder, configure the data mover options file for application protection:

- a) Open the data mover options file (`dsm.hostname_HV_DM.opt`) for editing.
For example, issue the following command:

```
notepad dsm.KINGSTON5_HV_DM.opt
```

- b) Add the `include.vmtsmvss guest_vm_name` option. The `guest_vm_name` parameter can contain wildcard characters.
For example:

```
include.vmtsmvss Kingston40
```

Or

```
include.vmtsmvss Kingston*
```

Alternatively, if you plan to manually preserve the SQL Server logs and restore SQL Server transactions to a specific checkpoint after the VM is restored, specify the following option to not truncate the SQL Server log:

```
include.vmtsmvss Kingston40 OPTions=KEEPSqllog
```

3. From the `baclient` folder, back up the guest VM by issuing the **dsmc backup vm** command:
For example:

```
dsmc backup vm Kingston40 -optfile=dsm.KINGSTON5_HV_DM.opt  
-asnode=KINGSTON5_HV_TGT
```

4. Verify the backup operation by running the **dsmc query vm** command. VM names are case-sensitive.
For example, issue the following command from the `baclient` folder:

```
dsmc query vm Kingston40 -optfile=dsm.KINGSTON5_HV_DM.opt  
-asnode=KINGSTON5_HV_TGT -detail
```

The output includes text that is similar to the following example (although your version of Microsoft SQL Server might be different):

```
Application protection type: TSM VSS  
Application(s) protected: Microsoft SQL Server 2017
```

5. From the `baclient` folder, use the **dsmc set access** command to allow the VSS node on the guest VM to access and restore the VM backup.
For example, issue a command that is similar to the following example:

```
dsmc set access backup -type=vm Kingston40 kingston40_vss  
-nodename=KINGSTON5_HV_TGT -optfile=dsm.KINGSTON5_HV_DM.opt
```

where `Kingston40` is the name of the guest VM that is shown in Microsoft Hyper-V Manager and `kingston40_vss` is an IBM Storage Protect node that is created during the installation of Data Protection for Microsoft SQL Server. This node name is in the file `C:\Program Files\Tivoli\TSM\baclient\dsm.opt` on the `Kingston40` VM.

Issue the following **dsmc query access** command to display the VM backups that the node has access to. For example:

```
dsmc query access -nodename=KINGSTON5_HV_TGT -optfile=dsm.KINGSTON5_HV_DM.opt
```

What to do next

When the backup operation is completed successfully, proceed to [“Step 4 \(Guest VM\): Restore a database”](#) on page 110.

Related reference

[“Backup VM”](#) on page 147

Use the **backup vm** command to back up Hyper-V virtual machines.

[“INCLUDE.VMTSMVSS”](#) on page 179

Related information

[Set Access](#)

[Set Password](#)

Step 4 (Guest VM): Restore a database

To verify that you correctly configured application protection, restore a database with Data Protection for Microsoft SQL Server.

Before you begin

Ensure that you completed the procedure in [“Step 3 \(Hyper-V host\): Configure Data Protection for Microsoft Hyper-V for application protection”](#) on page 107.

Ensure that the required services are running on the guest virtual machine (VM):

1. At the command prompt, issue the following command:

```
services.msc
```

2. Locate **IBM Storage Protect Recovery Agent Service** in the list of services and, if necessary, start the service.
3. Locate **Microsoft iSCSI Initiator Service** in the list of services. If necessary, change the startup type to **Automatic** and start the service.

Procedure

Complete the following steps on the guest VM:

1. Start the Data Protection for Microsoft SQL Server Management Console.
2. Select an **SQL Server** instance in the tree.
3. Navigate to the **Recover** tab, and click **Refresh**.
4. Select a database entry that has the **VMVSS** backup method.
5. In the **Action** pane, click **Restore to Alternate Location**. Specify a new database name and a new location where the database will be restored.
6. In the **Action** pane, click **Restore**.
7. When the restore operation is complete, review the database and any related tables.

What to do next

You can now manage backups and recover data if necessary. For more information, see:

- [“Managing backup operations” on page 112](#)
- [“Restoring data” on page 115](#)

If you change the name of the guest VM after you completed the configuration steps for application protection, you must reconfigure the software with the new VM name. For instructions, see [“Optional: Configuring application protection after a virtual machine name change” on page 110](#).

Optional: Configuring application protection after a virtual machine name change

If you changed the name of the guest virtual machine (VM) after you completed the application protection configuration, you must reconfigure Data Protection for Microsoft Hyper-V with the renamed guest VM.

Before you begin

Ensure that you have installed and configured the software to protect guest VMs that host Microsoft SQL Server.

About this task

Complete this task only if you changed the name of a guest VM that is protected by application protection.

Procedure

1. On the data mover on the Hyper-V host or cluster, issue the following command:

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

where:

new_vmguest_displayname

The name of the new guest VM that hosts Microsoft SQL Server. The name is the VM name that is displayed in Hyper-V Manager.

guest_admin_ID

The administrator ID for the new guest VM. The *guest_admin_ID* must be a Windows domain account or a local account. For example:

- For a domain account, use the *domain\username* format.
- For a local account, use the *username* format.

guest_admin_pw

The password for the administrator ID for the new guest VM.

2. In the data mover options file (*dsm.hostname_HV_DM.opt*), update the *include.vmtsmvss* option as follows:

```
include.vmtsmvss new_vmguest_displayname
```

where *new_vmguest_displayname* is the display name of the new guest VM in Hyper-V Manager. You can use wildcard characters.

3. From the *baclient* folder on the data mover, back up the new guest VM by using the Data Protection for Microsoft Hyper-V Management Console or by issuing the following command at the command prompt:

```
dsmc backup vm new_vmguest_displayname -optfile=dsm.hostname_HV_DM.opt  
-asnode=hostname_HV_TGT
```

4. From the *baclient* folder, use the **dsmc set access** command to allow the VSS node on the new guest VM to access and restore the VM backup.
For example, issue the following command:

```
dsmc set access backup -type=VM new_vmguest_displayname vss_requestor_node  
-optfile=dsm.hostname_HV_DM.opt
```

where:

new_vmguest_displayname

The display name of the new guest VM in Hyper-V Manager.

vss_requestor_node

The name of the VSS requester that was configured on Data Protection for Microsoft SQL Server.

hostname

The name of the Hyper-V host or cluster where Data Protection for Microsoft Hyper-V is installed.

5. Optional: Because the VSS requester node already has access to the VMs that were backed up under the old VM name, the Data Protection for Microsoft SQL Server Management Console shows the VMVSS databases that were backed up from both the old VM and the new VM.

If you do not want to access the VM backups with the old VM name, you must delete the access to the old VM backups. Issue the following command from the data mover on the Hyper-V host:

```
dsmc del access -optfile=dsm.hostname_HV_DM.opt -asnode=hostname_HV_TGT
```

An access list is displayed. Enter the index that corresponds to the item that you want to delete from the access list.

If you want to save space on the IBM Storage Protect server, you can delete the file space that contains the backup data for the old VM by issuing the **dsmc delete filespace** command.

Important: When you delete a file space, you delete all backup versions within that file space and you can no longer restore the data. Verify that the old VM backups are obsolete before you delete them.

Managing backup operations

After you configure the environment for to protect of Microsoft SQL Server data, you can schedule backups. You can set up schedules for a virtual machine (VM) backup operation and a Microsoft SQL Server log backup operation.

Scheduling virtual machine backups

To ensure that your data is protected regularly, schedule virtual machine (VM) backups.

Before you begin

By default, the maximum size that is allowed for a virtual hard disk (VHDX) in a backup operation is 2 TB. However, you can increase the maximum size to 8 TB by using the `vmmaxvirtualdisks` option. For more information, see [“Vmmaxvirtualdisks”](#) on page 204.

Procedure

1. Start the [Data Protection for Microsoft Hyper-V Management Console](#).
2. In the navigation pane, click a stand-alone host or cluster from the navigation pane.
3. In the **Actions** pane, click **Backup Management**.
4. Select a schedule in the **Backup Management** window and click **Assign Schedule**.
5. To close the window, click **Close**.
6. Verify that the source of the schedule includes the VMs that are hosting Microsoft SQL Server.
7. Verify that one of the following services is running:
 - If you are using a scheduler that is managed by a client acceptor (CAD), ensure that the client acceptor 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, 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 (VM) 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 VM.

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.

Procedure

1. From Microsoft Management Console (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 Microsoft Hyper-V are identified with the backup method `vmvss`.

Microsoft SQL Server logs that are backed up 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 (VM) 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 options file to specify the management class that is used for the VM backups.

Scheduled VM backups are associated with Data Protection for Microsoft Hyper-V.

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 VM, 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 Microsoft Hyper-V. 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 **REONLY** value for log backups to match the **RETEXTTRA** parameter for full database backups.

Verifying that Microsoft SQL Server volumes are not excluded in virtual machine backups

The volumes in Hyper-V virtual hard disks (VHDXs) must contain the Microsoft SQL Server databases that are not excluded from the Data Protection for Microsoft Hyper-V backup processing.

Before you begin

Ensure that the databases are not on any of the following types of disks:

- Physical disks
- Independent disks
- 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 Microsoft Hyper-V data mover that is used to back up the virtual machine (VM) do not inadvertently exclude VHDXs that are hosting volumes that contain Microsoft SQL Server files, file space, and database.

For example:

- `kingston40.vhdx` contains logical volume C:
 - `kingston40.vhdx` contains logical volumes E: and F:
 - The disk location (IDE controller number and device location) for `kingston40_1.vhdx` is IDE 1 0.
 - The disk location for `kingston40_2.vhdx` is IDE 1 1.
 - The Microsoft SQL Server database files to be backed up are on the E: and F: drive.
2. Verify that no statements exclude `kingston40_2.vhdx` from the VM backup by ensuring that the data mover does not contain the following or similar statements:

```
EXCLUDE.VMDISK KINGSTON40 "IDE 1 1"
EXCLUDE.VMDISK * "IDE 1 1"
```

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

```
INCLUDE.VMDISK KINGSTON40 "IDE 1 1"  
INCLUDE.VMDISK * "IDE 1 1"
```

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 VM disk from the command-line interface:

```
dsmc backup vm "KINGSTON40:-vmdisk=IDE 1 1" -asnode=KINGSTON5_HV_TGT
```

Related reference

[“Exclude.vmdisk” on page 172](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

[“Include.vmdisk” on page 176](#)

The `INCLUDE.VMDISK` option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

Restoring data

After you back up a virtual machine with the application protection feature enabled, you can recover a database in case the original is lost or damaged.

A recovery operation restores a full backup of the Microsoft SQL Server database from the Data Protection for Microsoft Hyper-V backup.

If you restore the entire virtual machine (VM), all Microsoft SQL Server databases on the VM are restored and recovered to the point of the VM 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

Complete the following steps in Windows **Services**.

1. In the **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 (VM) backup by using the Data Protection for Microsoft SQL Server graphical user interface.

Before you begin

Ensure that the Microsoft iSCSI Initiator Service is running before you restore any Microsoft SQL Server database with backup type "VMVSS". If the service is not running, start it. For instructions, see [“Starting the Microsoft iSCSI Initiator Service” on page 115](#).

Procedure

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

Restoring data by using the command-line interface

If you prefer, you can use the Data Protection for Microsoft SQL Server command-line interface to start a full Microsoft SQL Server database recovery from a virtual machine (VM).

Before you begin

Ensure that the Microsoft iSCSI Initiator Service is running before you restore any Microsoft SQL Server database with backup type "VMVSS". If the service is not running, start it. For instructions, see [“Starting the Microsoft iSCSI Initiator Service” on page 115](#).

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 6.0
(C) Copyright IBM Corporation 1997, 2018. All rights reserved.

Connecting to IBM Storage Protect Server as node 'KINGSTON40_SQL'...

Querying IBM Storage Protect Server for Backups ....

Backup Object Information
-----

SQL Server Name ..... SQL40
SQL Database Name ..... sql_db10
Backup Method ..... VMVSS
Backup Location ..... Srv
Backup Object Type ..... Full
Mount Points Root Directory .....
Backup Object State ..... Active
Backup Creation Date / Time ..... 07/12/2018 13:08:45
Backup Size ..... 17.00 MB
Backup Compressed ..... Yes
Backup Encryption Type ..... None
Backup Client-Deduplicated ..... Yes
Backup Supports Instant Restore ..... No
Database Object Name ..... 20180712130845
Assigned Management Class ..... STANDARD
Backup Modified .....
```

The operation completed successfully. (rc = 0)

2. To restore the database without applying transaction logs, issue the following command:

```
tdpsqlc restore databaseName /backupMethod=vmvss
```

The following example shows the output of the command when you specify the Microsoft SQL Server database called sql_db10.

```

tdpsqlc restore sql_db10 /backupmethod=vmvss /sqlserver=sql40
/fromsqlserver=sql40 /recovery=no

IBM Storage Protect for Databases:
Data Protection for Microsoft SQL Server
Version 8, Release 1, Level 6.0
(C) Copyright IBM Corporation 1997, 2018. All rights reserved.

Connecting to SQL Server, please wait...

Querying IBM Storage Protect Server for Backups ....

Connecting to IBM Storage Protect Server as node 'KINGSTON40_SQL'...
Connecting to Local DSM Agent 'SQL40'...
Using backup node 'KINGSTON40_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 log=* /sqlsever=sql40 /fromserver=sql40
/recovery=yes

```

You can also use the /stopat option to specify a more granular point in time.

```

tdpsqlc restore sql_db10 log=* /sqlserver=sql40
/fromsqlserver=sql40 /recovery=yes

IBM Storage Protect for Databases:
Data Protection for Microsoft SQL Server
Version 8, Release 1, Level 6.0
(C) Copyright IBM Corporation 1997, 2018. All rights reserved.

Connecting to SQL Server, please wait...
Starting Sql database restore...
Connecting to IBM Storage Protect Server as node 'KINGSTON40_SQL'...
Querying IBM Storage Protect server for a list
of database backups, please wait...

Beginning log restore of backup object sql_db10\20180712130845\000000DB0,
1 of 3, to database sql_db10 ...

Beginning log restore of backup object sql_db10\20180712130845\000000DB0,
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=20180712130845 /backupdestination=tsm  
/backupmethod=vmvss
```

Restoring Microsoft SQL Server log backups

After the full database is restored successfully, you can use Data Protection for Microsoft SQL Server to restore transaction logs.

Procedure

Complete the following steps on the guest VM:

1. Start the Microsoft Management Console (MMC) for Data Protection for Microsoft SQL Server.
2. Select a Microsoft SQL Server, and click the **Recover** tab.
3. Verify that the **AutoSelect** option is set to **False**.
4. Change the **RunRecovery** option to **True**.
5. Select the logs that you want to recover.
6. Click **Restore**.

Restoring relocated and deleted databases

The backup solution for restoring databases and log files that are relocated and deleted after a virtual machine (VM) backup requires Data Protection for Microsoft Hyper-V 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 Microsoft Hyper-V 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 Microsoft Hyper-V to back up the VM.
Consider the following example.

You back up VM kingston40 that includes Microsoft SQL Server database moose at 2:00 PM 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 PM 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 PM 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 PM

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 Microsoft Hyper-V 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 VM 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 (VM) backup. One procedure for checking for the existence of a full VM backup is to schedule the usage of a script.

The following sample script checks for the instance of a full backup and then runs the Microsoft SQL Server log backups if a full VM 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 6.0
(C) Copyright IBM Corporation 1997, 2018. 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: 20180703011509\000007CC
Backup of model completed successfully.
Beginning log backup for database sqldb test2, 2 of 2.
Full: 0 Read: 88832 Written: 88832 Rate: 132.44 Kb/Sec
Database Object Name: 20180703011511\000007CC
Backup of sqldb 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 on the IBM Storage Protect server to determine whether VM backups are successful.

IBM Storage Protect file space information

You might never need to know the file names or locations for your virtual machine (VM) files. However, if the underlying file structure interests you, Data Protection for Microsoft Hyper-V backups are stored under the node name of the Hyper-V target node (for example, KINGSTON5_HV_TGT) on the IBM Storage Protect server.

The following example shows the file space information for the virtual machine that is called kingston40.

```
Protect: ORION>q file KINGSTON5_HV_TGT f=d

Node Name: KINGSTON5_HV_TGT
Filespace Name: \VMFULL-kingston40
Hexadecimal Filespace Name:
FSID: 61
Collocation Group Name:
Platform: TDP Hyper-V
Filespace Type: API:TSMVM
Is Filespace Unicode?: No
Capacity: 0 KB
Pct Util: 0.0
Last Backup Start Date/Time: 03/13/2018 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:
```

Troubleshooting application protection of guest virtual machines

If Data Protection for Microsoft Hyper-V is configured for application protection of virtual machines (VMs) that host application data and you encounter a problem during VM backup operations, try to reproduce the problem in your environment.

Snapshots created during application protection backups that were interrupted cannot be deleted

If Data Protection for Microsoft Hyper-V is configured for application protection of VMs that host application data, you can run an application protection backup of a VM by issuing the `dsmc backup vm vmname` command. However, if you cancel the backup operation by pressing the **Ctrl + C** keys, the snapshot that is created by the backup operation is not removed automatically. Furthermore, the snapshot cannot be removed by using the Hyper-V Manager.

To resolve this problem, you must remove the snapshot manually by running the **Get-VMSnapshot** cmdlet with the **-SnapshotType Recovery** parameter, and then run the **Remove-VMSnapshot** cmdlet to remove the snapshot. For more information, see [“Snapshot management with Windows PowerShell”](#) on page 8.

Message ANS4063W is generated during an application protection backup of a Microsoft Exchange Server database

If you did not generate the credentials file on the guest VM and you back up the guest VM from the data mover, message ANS4063W 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.
Check health of application writers and databases.
```

To resolve this issue, complete the following steps:

1. In the guest VM, generate the credentials file in the guest by running the following command at a PowerShell command prompt and enter the domain user name (*domain name\user name*) and password when prompted:

```
Get-Credential | Export-Clixml -Path `
"C:\program files\Tivoli\TSM\baclient\dsmcreds.xml"
```

The domain user must have Exchange restore permission.

2. Verify the credentials by running the following commands in the guest VM from an Exchange Management Shell:

```
$cred = Import-CliXml -Path `
"C:\program files\Tivoli\TSM\baclient\dsmcreds.xml"
$Session = New-PSSession -Credential $cred -ConfigurationName `
Microsoft.Exchange -ConnectionUri `
http://<Exchange_server_name>/PowerShell?serializationLevel=Full `
-Authentication Kerberos
Import-PSSession -Session $Session
Get-MailboxDatabase -Server <Exchange_server_name>
```

3. On the data mover on the Hyper-V host or cluster, back up the guest VM by using Data Protection for Microsoft Hyper-V.

For instructions, see [“Running an ad hoc backup of a virtual machine” on page 77](#).

Related tasks

[“Troubleshooting VSS backup and restore operations on guest virtual machines” on page 122](#)

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

Troubleshooting VSS backup and restore operations on guest virtual machines

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

About this task

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

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 data mover 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 Linux system, the option file is `dsm.sys`. If the guest VM and data mover 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 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 Microsoft Exchange Server:

- The Exchange Server is not a member of a Database Availability Group (DAG).
- All Exchange 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 Server databases before you start the VM backup operation.

When Exchange Server DAG databases or Exchange Server 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 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 Microsoft Hyper-V backup with application protection of a VM might generate the following error in the `dsmererror.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 Microsoft Hyper-V 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 Microsoft Hyper-V resends the XML file that is identified in the error message in a new transaction.

PowerShell might run out of memory during database restore operations

When you restore Microsoft Exchange Server databases with the Data Protection for Microsoft Exchange Server Management Console, Windows PowerShell might fail with the following error:

```
APPCRASH
Not available
0
powershell.exe
10.0.14409.1005
584a185c
tsmapi64.dll
8.1.6.14
5af94075
c00000fd
00000000022df88
```

Exception code `0xc00000fd` indicates that a stack overflow exception occurred. To resolve this problem, increase the maximum amount of memory that is allocated for PowerShell by using the **MaxMemoryPerShell1MB** quota.

For details on how to change the value of the **MaxMemoryPerShell1MB** quota by using the Group Policy Editor (**gpedit.msc**) or PowerShell, see [https://msdn.microsoft.com/en-us/library/ee309367\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/ee309367(v=vs.85).aspx). The default value is 150, but a value of 1024 is preferred.

Chapter 7. Running instant access operations

You can run instant access operations to verify the integrity of a virtual machine (VM) backup or to clean up resources that are created by instant access operations.

About this task

You can verify that a previously backed up VM image can be used to successfully restore a system if the original VM is deleted or its disks and data are corrupted or otherwise unusable.

When you run an instant access operation, a temporary virtual machine is created. You can use this temporary VM for verification of the backup data, but the original VM is not restored.

To run instant access operations, you can use the data mover command line, the Data Protection for Microsoft Hyper-V Management Console, or Windows PowerShell cmdlets.

If you run instant access operations in a cluster environment, any cluster node can perform instant access operations by using its host iSCSI service, regardless of which host the VM is running on, and regardless of which host the instant access operation is started from.

Performance tip:

If you use instant access operations often, consider changing the location of the temporary write cache for the virtual volume. The default location is on the local hard disk (c:\ProgramData\Tivoli\TSM\RecoveryAgent\mount\). This location is used to temporarily store VM disk write changes from a VM running in instant access mode in the power-on state.

To help reduce the input/output (I/O) latency in the Hyper-V guest, use the IBM Storage Protect recovery agent GUI to specify the fastest disk on the Hyper-V host, such as a solid-state drive (SSD) or a SAN volume, for the temporary write cache. Any updates made to a VM running in instant access mode are only temporarily saved. When the instant access VM is cleaned up, the temporary write cache will not be saved to the IBM Storage Protect server. All files in the temporary write cache path will be removed. However, a folder that is named after the VM will remain.

For more information about changing the settings for the virtual volume write cache, see [“Configuring the IBM Storage Protect recovery agent GUI” on page 50](#).

Verifying the integrity of a VM backup from the command line

To verify the integrity of a backup image of a virtual machine (VM) without restoring the VM or disks to the Hyper-V host, run a full VM instant access operation on the data mover command line.

Before you begin

Review the information in [Chapter 7, “Running instant access operations,” on page 125](#).

Procedure

Complete the following steps on the data mover command line on the Hyper-V host or cluster:

1. Run a VM instant access operation by using the **restore vm** command.

Because the existing VM is running, you must provide a new name for a temporary VM by adding the new name to the **vmname** option. You must also add the **-VMRESToretype=INSTANTAccess** option to the command line to indicate that the operation is an instant access restore operation.

If you want the temporary VM to be automatically powered on, use the **-vmautostartvm=yes** option.

For example, you want to verify that the backup image of the VM named Orion can be used for restore operations. Enter the following command to prepare a VM named "Orion_verify" for instant access. You can use this VM to verify that the backup image can be restored.

```
dsmc restore vm Orion -vmname=Orion_verify -vmrestoretype=instantaccess
```

Important: If the VM that you are restoring is using a static IP address, a network IP address conflict will occur between the original VM and the temporary VM. To avoid this conflict, change the IP address of the temporary instant access VM after the VM is created.

2. Optional: Select an inactive or active VM backup from a particular date or time, list the versions of the VM backup by using the **inactive** or **pick** options, or the **pittime** or **pitdate** options.

For example, to display a list of backup versions of the Orion VM, issue the following command:

```
dsmc restore vm Orion -vmrestoretype=instantaccess -pick
```

Temporary data that is created by the instant access operation is stored in the cache of the IBM Storage Protect recovery agent.

3. To verify the integrity of the disks and data, use a utility such as **chkdsk**, or a utility or application of your choice, to verify the virtual disks and data. If the temporary VM passes the integrity checks, you can remove the temporary resources that were created to support the instant access restore operation.

Example

Run an instant access operation for a VM backup named "Ritza2-VM" and restore the VM temporarily to "Ritza2-VM_Verify". In this example, the target node name is "RITZA_HV_TGT", and the data mover node name is "RITZA_HV_DM":


```
dsmc restore vm "Ritza2-VM" -vmname="Ritza2-VM_Verify" -vmrestoretype=instantaccess
-optfile="dsm.RITZA_HV_DM.opt" -asnode="RITZA_HV_TGT"
```

IBM Storage Protect

Command Line Backup-Archive Client Interface

Client Version 8, Release 1, Level 7.0

Client date/time: 07/19/2018 13:06:23

(c) Copyright by IBM Corporation and other(s) 1990, 2018. All Rights Reserved.

Node Name: RITZA_HV_DM

Session established with server ASANA2: Linux/x86_64

Server Version 7, Release 1, Level 9.000

Server date/time: 07/19/2018 13:04:32 Last access: 07/19/2018 13:03:56

Accessing as node: RITZA_HV_TGT

Restore function invoked.

Restore VM command started. Total number of virtual machines to process: 1

Restore of Virtual Machine 'Ritza2-VM' started

Starting instant access of VMware Virtual Machine 'Ritza2-VM' target node name='RITZA_HV_TGT', data mover node name='RITZA_HV_DM'

Starting Instant VM Access process

Restoring VM configuration information for 'Ritza2-VM'

Checking the state of the hypervisor host 'RITZA'...

Checking Hyper-V target path 'C:\ProgramData\Microsoft\Windows\Hyper-V' from host 'RITZA'

Gathering iSCSI information from Hyper-V host 'RITZA' ...

The Windows iSCSI Initiator name is "iqn.1991-05.com.microsoft:ritza.tsmdev.local".

Mounting a snapshot of VM 'Ritza2-VM' from disc 'Hard Disk 1'. Please wait...

Connecting the iSCSI targets with the data mover...

Creating a virtual machine on the hypervisor host...

Successful Full VM restore of Virtual Machine 'Ritza2-VM'

ANS2408I The virtual machine named 'Ritza2-VM_Verify' is ready for Instant Access

Successful instant access of VMware Virtual Machine 'Ritza2-VM' target node name='RITZA_HV_TGT', data mover node name='RITZA_HV_DM'

Total number of objects restored:	1
Total number of objects failed:	0
Total number of bytes transferred:	0 B
Data transfer time:	0.00 sec
Network data transfer rate:	0.00 KB/sec
Aggregate data transfer rate:	0.00 KB/sec
Elapsed processing time:	00:01:03

Related tasks

[“Freeing instant access resources by using the command line” on page 128](#)

You can free the resources that were created by instant access operations by running a cleanup operation.

[“Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console” on page 130](#)

For ease of use, you can run instant access operations from the Data Protection for Microsoft Hyper-V Management Console.

Related reference

[“Vmrestoretype” on page 209](#)

Use the **vmrestoretype** option with the **query VM** or **restore VM** commands to specify the type of restore operation to perform or query.

[“Vmautostartvm” on page 189](#)

Use the `vmautostartvm` option to specify whether the virtual machine (VM) that is created during instant access processing is automatically powered on. This option is valid only when used with the **restore VM** `vmrestoretype=instantaccess` command.

[“PowerShell cmdlets for Data Protection for Microsoft Hyper-V” on page 135](#)

Review the Data Protection for Microsoft Hyper-V cmdlets that you can use to protect your virtual machine (VM) data.

Related information

[“Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143](#)

Freeing instant access resources by using the command line

You can free the resources that were created by instant access operations by running a cleanup operation.

About this task

You must determine whether any temporary instant access virtual machines (VMs) exist, and then run a clean-up operation on the data mover command line.

Procedure

1. Run the **query vm** command with the `-VMRESToretype=INSTANTAccess` option. This command displays all temporary VMs that are running in instant access mode:

```
dsmc query vm * -VMREST=INSTANTA
```

You can add the **-Detail** option to the **query vm** command to display more information about each of the temporary VMs. For example, issue the following command:

```
dsmc query vm vmname -VMREST=INSTANTA -detail
```

where *vmname* specifies the name of a temporary VM.

2. Remove the resources that were created by the instant access operation.

For example, to remove the resources that were created for a temporary VM named "Orion_verify", issue the following command:

```
dsmc restore vm Orion -vmname=Orion_verify -VMRESToretype=VMCleanup
```

The **-VMRESToretype=VMCleanup** option deletes the temporary VM from the Hyper-V host, unmounts any iSCSI mounts that were mounted, and clears the iSCSI device list from the Hyper-V host. Deleting the VM also deletes all temporary generated data.

Example

Run an instant access cleanup operation to remove resources that were created by the instant access operation.

The following example shows the command that is used to clean up the temporary VM "kingston30-avi_Verify". In this case, the VM name is "kingston30-avi", the target node name is "KINGSTON10_HV_TGT", and the data mover node name is "KINGSTON10_HV_DM".

```
dsmc restore vm "kingston30-avi" -vmname="kingston30-avi_Verify"  
-vmrestoretype=vmcleanup  
-optfile="dsm.KINGSTON10_HV_DM.opt" -asnode="KINGSTON10_HV_TGT"
```

IBM Storage Protect

Command Line Backup-Archive Client Interface

Client Version 8, Release 1, Level 7.0

Client date/time: 07/25/2018 09:13:01

(c) Copyright by IBM Corporation and other(s) 1990, 2018. All Rights Reserved.

Node Name: KINGSTON10_HV_DM

Session established with server ASANA2: Linux/x86_64

Server Version 7, Release 1, Level 9.000

Server date/time: 07/25/2018 09:11:54 Last access: 07/25/2018 09:09:33

Accessing as node: KINGSTON10_HV_TGT

Restore function invoked.

Restore VM command started. Total number of virtual machines to process: 1

Restore of Virtual Machine 'kingston30-avi' started

Starting cleanup of VMware Virtual Machine 'kingston30-avi_Verify' from snapshot of 'kingston30-avi' target node name='KINGSTON10_HV_TGT', data mover node name='KINGSTON10_HV_DM'

Starting Instant VM Access cleanup process

Restoring VM configuration information for 'kingston30-avi'

Checking the state of the hypervisor host 'KINGSTON10'...

Powering off the virtual machine ...

Removing the virtual machine from the Hyper-V host...

Disconnecting iSCSI targets from the data mover...

Dismounting a snapshot of VM 'kingston30-avi' from target 'kingston30-avi9verify-hard8disk31'.

Please wait...

Cleanup operation of instant access VM 'kingston30-avi_Verify' completed successfully

Successful cleanup of VMware Virtual Machine 'kingston30-avi_Verify' from snapshot of 'kingston30-avi' target node name='KINGSTON10_HV_TGT', data mover node name='KINGSTON10_HV_DM'

Total number of objects restored:	1
Total number of objects failed:	0
Total number of bytes transferred:	0 B
Data transfer time:	0.00 sec
Network data transfer rate:	0.00 KB/sec
Aggregate data transfer rate:	0.00 KB/sec
Elapsed processing time:	00:00:22

What to do next

You can also use the **-restoreType VmCleanup** parameter with the **Restore-DpHvVm** cmdlet to free instant access resources. For more information, see [“Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143](#).

Related reference

[“Vmrestoretype” on page 209](#)

Use the `vmrestoretype` option with the **query VM** or **restore VM** commands to specify the type of restore operation to perform or query.

[“Query VM” on page 154](#)

Use the **query VM** command to list and verify the successful backups of virtual machines (VMs).

Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console

For ease of use, you can run instant access operations from the Data Protection for Microsoft Hyper-V Management Console.

Before you begin

Review the information in [Chapter 7, “Running instant access operations,” on page 125](#).

About this task

When you run an instant access operation, a temporary virtual machine (VM) is created. You can use this temporary VM for verification of the backup data, but the original VM is not restored.

Procedure

Complete the following steps on the Hyper-V host or cluster to create a temporary instant access VM:

1. Start the Data Protection for Microsoft Hyper-V Management Console.
2. In the navigation pane, click a host node or the cluster node.
3. Select a VM from the VM table in the **Virtual Machines** view.
For example, click a VM whose data protection status is **Normal**.

Tip: You can also run instant access operations on VMs with the **Deleted** status.

4. In the **Actions** pane, click **Restore**.
5. Complete the following pages in the **Restore** wizard as applicable. The pages that are provided depend on the options that you select in the wizard.

Wizard page	Action
Before you begin	Click Next to start the wizard.
Select restore point	<p>The highlighted dates on the calendar contain restore points. Restore points are VM backups that are available for restore operations. Some VMs have more than one restore point per day.</p> <p>Select a date and a restore point from the Available restore points list. The size of the VM is listed next to an available restore point. The VM will be restored to the state in which it existed when it was backed up.</p>
Select options	<p>To create an instant access VM, complete the following steps:</p> <ol style="list-style-type: none">a. Click one of the following options:<ul style="list-style-type: none">Create an instant access virtual machine Creates a temporary VM for use with instant access operations. Does not power on the VM.Create and start an instant access virtual machine Creates a temporary VM for use with instant access operations and powers on the VM automatically.b. In the Virtual machine name field, enter a new name for the temporary instant access VM.

Wizard page	Action
	c. For cluster configurations, in the Restore virtual machine to field, select a host where you want to create the instant access VM. You can select the same host that the original VM is on or select a different host in the cluster.
Summary	Review the options that you selected in the wizard. Click Next to start the restore operation.
Results	Click Finish to close the wizard.

Important: If the VM that you are restoring is using a static IP address, a network IP address conflict will occur between the original VM and the temporary VM. To avoid this conflict, change the IP address of the temporary instant access VM after the VM is created.

Results

The restore operation that you started is displayed in the task list in the results pane of the management console.

After the restore operation is completed, the instant access VM is displayed in the **Virtual Machines** view or **Schedule History** view. The VM is shown with the **Instant Access** type in the **VM Type** column.

Temporary data that is created by the instant access operation is stored in the cache of the IBM Storage Protect recovery agent.

What to do next

To verify the integrity of the disks and data, use a utility such as **chkdsk**, or a utility or application of your choice, to verify the virtual disks and data. If the temporary VM passes the integrity checks, you can remove the temporary resources that were created to support the instant access restore operation.

To remove the instant access VM after you finish verifying the backup data, follow the instructions in [“Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console” on page 131](#)

Related tasks

[“Verifying the integrity of a VM backup from the command line” on page 125](#)

To verify the integrity of a backup image of a virtual machine (VM) without restoring the VM or disks to the Hyper-V host, run a full VM instant access operation on the data mover command line.

Related information

[“Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143](#)

Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console

When you no longer need an instant access virtual machine (VM), you can use the Data Protection for Microsoft Hyper-V Management Console to remove the VM.

About this task

Removing an instant access VM deletes the temporary VM from the Hyper-V host, unmounts any iSCSI mounts that were mounted, and clears the iSCSI device list from the Hyper-V host. Deleting the VM also deletes all temporarily generated data.

Procedure

To remove an instant access VM, complete the following steps on the Hyper-V host:

1. [Start the Data Protection for Microsoft Hyper-V Management Console.](#)

2. In the navigation pane, click a host in the cluster or host view.
3. In the VM table in the **Virtual Machines** view, select the instant access VM that you want to remove. The VMs that are in instant access mode are identified by the **Instant Access** entry in the **Status** column.
4. In the **Actions** pane, click **Delete Instant Access VM**.

When prompted, confirm that you want to remove the instant access VM.

Related tasks

[“Verifying the integrity of a VM backup from the command line” on page 125](#)

To verify the integrity of a backup image of a virtual machine (VM) without restoring the VM or disks to the Hyper-V host, run a full VM instant access operation on the data mover command line.

Related information

[“Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143](#)

Restrictions on instant access VMs after a Hyper-V host is restarted

When a Hyper-V host is restarted, a restriction exists for virtual machines (VMs) that were created by instant access operations.

By default, when a Hyper-V host is restarted, the VMs in the Hyper-V host will be changed to the saved state and will start automatically after the restart. This action does not apply to VMs that are created by an instant access operation. After a restart, the disks that were provided by the recovery agent no longer exist. To regain access to the instant access VMs, you must clean up and rerun the instant access operation.

Related tasks

[“Removing an instant access VM from the Data Protection for Microsoft Hyper-V Management Console” on page 131](#)

When you no longer need an instant access virtual machine (VM), you can use the Data Protection for Microsoft Hyper-V Management Console to remove the VM.

[“Running instant access operations from the Data Protection for Microsoft Hyper-V Management Console” on page 130](#)

For ease of use, you can run instant access operations from the Data Protection for Microsoft Hyper-V Management Console.

[“Freeing instant access resources by using the command line” on page 128](#)

You can free the resources that were created by instant access operations by running a cleanup operation.

[“Verifying the integrity of a VM backup from the command line” on page 125](#)

To verify the integrity of a backup image of a virtual machine (VM) without restoring the VM or disks to the Hyper-V host, run a full VM instant access operation on the data mover command line.

Chapter 8. Protecting virtual machines by using Windows PowerShell cmdlets

You can run Data Protection for Microsoft Hyper-V operations by using Microsoft Windows PowerShell cmdlets (Version 3.0 or later).

About this task

Information about preparing to use PowerShell cmdlets, the list of available cmdlets, and common tasks that use these cmdlets are provided.

Restriction: The PowerShell cmdlets interact with the Data Protection for Microsoft Hyper-V REST API to protect your virtual machines. You cannot interact with the REST API directly. You must use the provided PowerShell cmdlets to run Data Protection for Microsoft Hyper-V operations.

Preparing to use PowerShell cmdlets with Data Protection for Microsoft Hyper-V

Data Protection for Microsoft Hyper-V includes a set of Windows PowerShell cmdlets to help you manage Data Protection for Microsoft Hyper-V operations in your environment.

Before you begin

Ensure that Microsoft Windows PowerShell 3 or later is available on the system where Data Protection for Microsoft Hyper-V is installed. To view which version of PowerShell is installed, enter the following command in a PowerShell session:

```
PS C:\> $PSVersionTable.PSVersion
```

The number in the **Major** column is the PowerShell version.

About this task

You can run the cmdlets interactively at the PowerShell command line or include them in scripts that can help you automate Data Protection for Microsoft Hyper-V operations.

You must complete the following steps before you use the cmdlets.

Procedure

1. Start a Microsoft Windows PowerShell or Microsoft Windows PowerShell ISE session with administrator authority:
 - a) Click **Start > All Programs > Accessories > Windows PowerShell**.
 - b) Right-click **Windows PowerShell** and click **Run as administrator**.
2. Verify that the 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
```

Tip: The **Set-ExecutionPolicy** command must be run only once.

3. To make the cmdlets available, import the Data Protection for Microsoft Hyper-V PowerShell module:

```
PS C:\> Import-Module "C:\Program
Files\IBM\StorageProtect\DPHyperV\dphvModule.dll"
```

4. Authenticate to Data Protection for Microsoft Hyper-V by using the session cmdlet:

```
$cred = Get-Credential -UserName user_name -message "credential"
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
```

where:

user_name

Specifies the account that you use to log in to the Windows system where Data Protection for Microsoft Hyper-V is installed.

computer_name

Specifies the name of the server where Data Protection for Microsoft Hyper-V is installed.

5. If the security certificate that is associated with the host you are connecting to is not recognized or is not on the local server (where the PowerShell cmdlets are installed), the session cmdlet fails. You must rerun the **New-DpHvSession** cmdlet with either the **-Force** parameter to ignore the certificate or the **-CertificatePrompt** parameter to display a prompt for installing a new certificate. For Windows Core OS with or without Features on Demand installed, use the **-Force** parameter. The **-CertificatePrompt** parameter is not supported for installing a new certificate.

For example, run the following session cmdlet:

```
$cred = Get-Credential -UserName user_name -message "credential"
$session = New-DpHvSession -ComputerName computer_name -Credential $cred `
-CertificatePrompt
```

When you are prompted, complete the following steps for a stand-alone host or for each host in a cluster:

- a) In the **Securing connection to <host name>** window, click **View certificate**.

If you select any other options, such as **Yes** to ignore the certificate warning for the current session, **No** to stop the connection, or **Don't ask me again for connection to this computer** to ignore all future certificate warnings, you will not be able to connect to Data Protection for Microsoft Hyper-V.

- b) In the **General** tab of the **Certificate** window, click **Install Certificate**.
- c) In the welcome page of the **Certificate Import Wizard** window, select a store location (**Current User** or **Local Machine**) and click **Next**.
- d) In the **Certificate Store** page, click **Place all certificates in the following store** and click **Browse**.
- e) In the **Select Certificate Store** window, select **Trusted Root Certification Authorities** and click **OK**.
- f) Click **Next** in the **Certificate Store** page.
- g) Review the selections in the **Completing the Certificate Import Wizard** page and click **Finish**.
- h) In the **Security Warning** window, click **Yes** to install the certificate.
- i) Click **OK** in the confirmation window.

If you reject the certificate, you will not be able to connect to Data Protection for Microsoft Hyper-V unless you use the **-Force** parameter.

6. Review the list of available cmdlets in [“PowerShell cmdlets for Data Protection for Microsoft Hyper-V” on page 135](#).
7. Optional: Review the online help for each cmdlet. For more information, see [“Getting help information for PowerShell cmdlets” on page 138](#).

What to do next

For information about creating, running, monitoring, and troubleshooting scripts with cmdlets, see Windows PowerShell 3.0 or later documentation. For more information about Windows PowerShell

cmdlets, consistent naming patterns, parameters, arguments, and syntax, see [Microsoft TechNet: Getting Started with Windows PowerShell](#).

PowerShell cmdlets for Data Protection for Microsoft Hyper-V

Review the Data Protection for Microsoft Hyper-V cmdlets that you can use to protect your virtual machine (VM) data.

The following table identifies the cmdlets that are available for Data Protection for Microsoft Hyper-V.

Table 14. PowerShell cmdlets for Data Protection for Microsoft Hyper-V	
Cmdlet name	Description
Backup-DpHvVm	Back up a Hyper-V VM. Related command: dsmc backup vm
Get-DpHvBackup	Show information about a VM backup. Related command: dsmc query vm
Get-DpHvBackupSchedule	Show a list of eligible backup schedules for the VMs in the Hyper-V host or cluster. An eligible schedule must be defined by the IBM Storage Protect server administrator, and it must be defined for a domain that is targeted for Hyper-V VMs. The schedule definition must include the following parameters and options: <ul style="list-style-type: none">• The <code>-domain.vmfull="all-vm"</code> option must be specified in the option string.• The schedule must contain the <code>ACTION=BACKUP</code> and <code>SUBACTION=VM</code> parameters.
Get-DpHvHostConfiguration	Show the configuration information for the Hyper-V host from the IBM Storage Protect server.
Get-DpHvIaVm	Show the list of VMs that are in instant access mode. Attributes about the VMs are also displayed. Instant access VMs are temporary VMs that are created for instant access operations. For more information, see “Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143.
Get-DpHvLastSuccessfulBackup	Show information about recent VM backup operations that ran in a host or cluster. The following information is returned: the at-risk status of the VM backup, the backup date, the duration of the backup, the amount of data that was transmitted, the type of backup, the host that the VM belongs to (for clusters), and the name of the schedule that ran.
Get-DpHvLog	Return the error log entries for a node. You can save the log entries to a file and specify a file name.
Get-DpHVLogSection	Return a section of the error log when you specify the starting line number and how many lines to display.

Table 14. PowerShell cmdlets for Data Protection for Microsoft Hyper-V (continued)	
Cmdlet name	Description
Get-DpHvPolicyDomain	Show a list of policy domains on the IBM Storage Protect server. Related command: dsmdmc q domain
Get-DpHvRCInfo	Show information about a return code.
Get-DpHvScheduleHistory	Show a list of schedules that ran. Each schedule that is returned can contain the time that the schedule started to run, the name of the schedule, the status of the schedule, the number of VMs that were backed up or failed to back up, and the duration of the schedule run.
Get-DpHvScheduleHistoryDetail	Show information for each VM that was backed up in a schedule run. Each backup task that is returned can contain the name of the VM, the status of the backup, the start time of the backup, and error codes for backups that failed.
Get-DpHvTask	Show general information about completed and running tasks.
Get-DpHvVvm	Show information about the VM inventory on the Hyper-V host. Related command: dsmc show vm
Get-DpHvVMAtRisk	Show the current at-risk policy for the VM. The at-risk policy determines when a VM backup is shown as at-risk if a backup operation did not occur in a specified time interval. The at-risk policy consists of an at-risk type. The at-risk type is a number and can be 0 (BYPASS), 1 (CUSTOM), or 2 (DEFAULT). Custom at-risk types have an at-risk interval in hours.
Get-DpHvVMBackupHistory	Show the backup history for a VM from the summary extended table on the IBM Storage Protect server. Each backup task that is returned can contain the last run time of a backup, the status of the backup, the duration of the backup, the amount of data that was transmitted, the host that the VM belongs to (for clusters), and any error codes that were returned.
Get-DpHvVmBackupTaskHistory	Show the VM backup task history that is stored locally on Data Protection for Microsoft Hyper-V.
Get-DpHvVmRestoreTaskHistory	Show the VM restore task history that is stored locally on Data Protection for Microsoft Hyper-V.

Table 14. PowerShell cmdlets for Data Protection for Microsoft Hyper-V (continued)	
Cmdlet name	Description
New-DpHvFsInfo	Create an FsInfo object for use with the Set-DpHvVmAtRisk cmdlet. The FsInfo object specifies the IBM Storage Protect file space ID and the name of the VM for which to set the at-risk policy.
New-DpHvNodeInfo	Create a NodeInfo object for use with the Set-DpHvHostConfiguration cmdlet. The NodeInfo object specifies the node name and node type on the Hyper-V host.
New-DpHvSession	Authenticate to Data Protection for Microsoft Hyper-V and start a PowerShell cmdlet session.
Receive-DpHvTask	Monitor the backup or restore task.
Remove-DpHvSession	Close the session with Data Protection for Microsoft Hyper-V.
Restore-DpHvVm	Restore a Hyper-V VM. You can also verify whether a VM backup can be restored by running instant access operations. For more information, see “Example 14: Verify the integrity of a VM backup by running cmdlets” on page 143. Related command: dsmc restore vm
Set-DpHvBackupSchedule	Associate a schedule with a data mover on a Hyper-V host or cluster. You can associate a node with a schedule, remove a node from the schedule, and associate a target node with a schedule.
Set-DpHvHostConfiguration	Configure the Hyper-V host for Data Protection for Microsoft Hyper-V operations.
Set-DpHvHttpsPort	Set the HTTPS port that is used for the IBM Storage Protect web server.
Set-DpHvMmcLoginPreferences	Set the preferences for the Data Protection for Microsoft Hyper-V Management Console.
Set-DpHvVmAtRisk	Set the at-risk policy for one or more VMs. The at-risk policy determines when a VM backup is shown as at-risk if a backup operation did not occur in a specified time interval. The at-risk policy consists of an at-risk type. The at-risk type is a number and can be 0 (BYPASS), 1 (CUSTOM), or 2 (DEFAULT). Custom at-risk types have an at-risk interval in hours.
Set-ServerConnection	Store IBM Storage Protect server connection information on the Hyper-V host and verify the connection to the server.
Show-DpHvHttpsPort	Show the port information for the IBM Storage Protect web server.

Table 14. PowerShell cmdlets for Data Protection for Microsoft Hyper-V (continued)	
Cmdlet name	Description
Show-DpHvMmcLoginPreferences	Show the preferences for the Data Protection for Microsoft Hyper-V Management Console.
Stop-DpHvTask	Cancel a backup or restore task.
Test-DpHvConfiguration	Verify the configuration for Data Protection for Microsoft Hyper-V.
Test-DomainCredentials	Verify the credentials for the Windows domain user.

For the list of common tasks for the cmdlets, see [“Data Protection for Microsoft Hyper-V cmdlet examples”](#) on page 138.

Getting help information for PowerShell cmdlets

Online help is provided for the PowerShell cmdlets. To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name:

```
Get-Help cmdlet_name
```

For example:

```
Get-Help Backup-DpHvVm
```

The following examples are for the **Backup-DpHvVm** cmdlet. To see cmdlet examples, enter:

```
Get-Help Backup-DpHvVm -examples
```

For detailed information, enter:

```
Get-Help Backup-DpHvVm -detailed
```

For technical information, enter:

```
Get-Help Backup-DpHvVm -full
```

For online product information, enter:

```
Get-Help Backup-DpHvVm -online
```

For information about a specific parameter, such as the **IFINCREMENTAL** parameter, enter:

```
help Backup-DpHvVm -Parameter IFINCREMENTAL
```

To show the help in a separate window, include the **-ShowWindow** parameter with the **help** command.

Data Protection for Microsoft Hyper-V cmdlet examples

Examples of Data Protection for Microsoft Hyper-V cmdlets are provided to help you protect your Hyper-V virtual machines (VMs).

Before you use the cmdlets, ensure that you complete the steps in [“Preparing to use PowerShell cmdlets with Data Protection for Microsoft Hyper-V”](#) on page 133.

Examples are provided for commonly used Data Protection for Microsoft Hyper-V tasks.

Tips:

- Each cmdlet provides parameters. To view the parameters, issue the following **help** command:

```
help cmdlet_name -ShowWindow
```

- Online help is available for the cmdlets. For more information, see [“Getting help information for PowerShell cmdlets”](#) on page 138.
- Unless stated otherwise in the examples, *user_name* specifies the account that you use to log in to the computer where Data Protection for Microsoft Hyper-V is installed. *computer_name* specifies the server where Data Protection for Microsoft Hyper-V is installed.

Examples

[“Example 1: Back up one or more VMs”](#) on page 139

[“Example 2: Query a VM backup”](#) on page 140

[“Example 3: Verify whether a Hyper-V host is configured for Data Protection for Microsoft Hyper-V operations”](#) on page 140

[“Example 4: Store IBM Storage Protect server connection information on the Hyper-V host and verify the connection”](#) on page 140

[“Example 5: Display the policy information on the IBM Storage Protect server”](#) on page 140

[“Example 6: Configure a Hyper-V host for Data Protection for Microsoft Hyper-V operations”](#) on page 140

[“Example 7: Show the VM inventory on the Hyper-V host”](#) on page 140

[“Example 8: Show the backup status of VMs on a host or cluster”](#) on page 141

[“Example 9: Set the at-risk policy for a VM”](#) on page 141

[“Example 10: Show the history of schedule runs”](#) on page 141

[“Example 11: Associate a schedule with a data mover on a host or cluster”](#) on page 142

[“Example 12: Restore one or more VMs”](#) on page 142

[“Example 13: Verify the configuration of Data Protection for Microsoft Hyper-V”](#) on page 143

[“Example 14: Verify the integrity of a VM backup by running cmdlets”](#) on page 143

Example 1: Back up one or more VMs

Run an incremental-forever incremental backup of one or more VMs.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
$vmList = @("vm1", "vm2")
$task = Backup-DpHvVm -Session $session -VmName $vmList -mode IFINCREMENTAL
$taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
while ("running" -eq $taskInfo.taskState) {
    start-sleep -seconds 30
    $taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
    if ($taskInfo.hasMoreData) {
        $results = Receive-DpHvTask -Session $session -TaskId $task.taskId
        write-verbose -verbose ("Started {0} Duration {1:g} Transferred `
            {2:N2} MB" -f $results.startTime, ((Get-Date) - $results.startTime), `
            ($results.totalBytesTransferred/1MB))
    }
}

$results = Receive-DpHvTask -Session $session -TaskId $task.taskId
$results

Remove-DpHvSession -Session $session
```

This example starts a PowerShell cmdlet session with Data Protection for Microsoft Hyper-V, backs up the VMs, queries the VM backup, monitors the backup job, and ends the session when the backup is completed.

Example 2: Query a VM backup

Query the IBM Storage Protect server file space and show general information about all VM backups.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
$bks = Get-DpHvBackup -Session $session
$bks
Remove-DpHvSession -Session $session
```

Example 3: Verify whether a Hyper-V host is configured for Data Protection for Microsoft Hyper-V operations

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
Get-DpHvHostConfiguration -Session $session
Remove-DpHvSession -Session $session
```

Example 4: Store IBM Storage Protect server connection information on the Hyper-V host and verify the connection

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
Set-ServerConnection -Session $session -ServerName server_name -Admin `
    admin_name -AdminPwd admin_password -ServerSSLPort port
Remove-DpHvSession -Session $session
```

Example 5: Display the policy information on the IBM Storage Protect server

Display information such as the domain name, default management class, description, and the duration of backup and archive retention:

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
Get-DpHvPolicyDomain -Session $session
Remove-DpHvSession -Session $session
```

Example 6: Configure a Hyper-V host for Data Protection for Microsoft Hyper-V operations

The following example configures a Hyper-V host by completing the following tasks:

- Register the target node (cluster node).
- Register the data mover node and configure it for backup operations (configure the options file, and create the client acceptor and scheduler services).
- Configure the file restore environment if requested (register the Windows and Linux mount proxy nodes, and create the options file and client acceptor services). If the file restore feature is enabled, the file restore credential must be the domain user and password.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
Set-ServerConnection -Session $session -ServerName server_name -Admin `
    admin_name -AdminPwd admin_password -ServerSSLPort port
$nodesList = @(New-DpHvNodeInfo -NodeName node_name -NodeType node_type)
Set-DpHvHostConfiguration -Session $session -PolicyDomain policy_domain_name `
    -RegisterTargetNode -TargetNode target_node -NodeList $nodesList -EnableFR `
    -FRDomainUser domain_name\user_name -FRDomainPwd password
Remove-DpHvSession -Session $session
```

Example 7: Show the VM inventory on the Hyper-V host

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
$vms = Get-DpHvVm -Session $session
```

```
$vms
Remove-DpHvSession -Session $session
```

Example 8: Show the backup status of VMs on a host or cluster

The following example returns information about the last VM backups on a host or cluster.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred
$lastBackups = Get-DpHvLastSuccessfulBackup -Session $session
$vmName = $lastBackups | select -first 1 -ExpandProperty name
$vmBackupHistory = Get-DpHvVMBackupHistory -Session $session -vmName $vmName
$vmBackupHistory
Remove-DpHvSession -Session $session
```

Example 9: Set the at-risk policy for a VM

The at-risk policy determines that a VM is at risk of being unprotected if a scheduled backup operation did not occur within a specified time interval.

The first half of the following example displays at-risk information for all VMs that have been backed up. The second half of the example updates the at-risk value for all VMs that begin with "SQL" to 12 hours.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

$lastBackups = Get-DpHvLastSuccessfulBackup -Session $session

# 1 - display the current at risk value for all vms

$i = 0
$atRiskList = @()
foreach ($backup in $lastBackups) {
    $activity = "Checking at risk value for {0}" -f $backup.name
    Write-Progress -activity $activity -status "Progress:" -percentcomplete `
        ($i++/$lastBackups.count*100)
    $atRisk = Get-DpHvVmAtRisk -session $session -VmName $backup.name
    $atRiskList += [pscustomobject]@{VM=$backup.name;AtRiskType=`
        $atRisk.AtRiskType;AtRiskInterval=$atRisk.AtRiskInterval}
}
$atRiskList | Out-GridView -Title "VM Risk Status" -PassThru

# 2 - set the at-risk value for all VMS that begin with SQL to a custom interval
# of 12 hours

$sqlVms = $lastBackups | where name -like "sql*"
$fsList = @()
foreach ($vm in $sqlVms) {
    $fsList += New-DpHvFsInfo -vmName $vm.Name -fsId $vm.FileSpaceId
}
Set-DpHvVmAtRisk -session $session -AtRiskType CUSTOM -AtRiskInterval 12 `
    -FsList $fsList

Remove-DpHvSession -Session $sess
```

Example 10: Show the history of schedule runs

The following example displays a summary of scheduled activity followed by the details of the most current scheduled activity.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

$schedHistory = Get-DpHvScheduleHistory -Session $session
$sh = $schedHistory | Sort-Object actualstarttime -Descending | Select-Object `
    -First 1
$schedHistoryDetail = Get-DpHvScheduleHistoryDetail -Session $session -ScheduleName
    $sh.Name -StartTime $sh.ActualStartTime -EndTime $sh.EndTime -NodeList `
    $sh.NodeList

#"Schedule History Summary"
$schedHistory |
    select actualstarttime,name,status,vmsucceeded,vmfailures,duration,nodelist | `
```

```

        sort actualstarttime -desc | ft -AutoSize

#"Details of most recent scheduled activity"
$schedHistoryDetail |
    select starttime,damamover,targetnode,name,status,duration,datatransmitted,`
    backuptype| ft -AutoSize

Remove-DpHvSession -Session $session

```

Example 11: Associate a schedule with a data mover on a host or cluster

You can verify a schedule association by running the QUERY ASSOCIATION command on the IBM Storage Protect server.

```

$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

# Get a list of schedules from the IBM Storage Protect server
$scheduleList = Get-DpHvBackupSchedule -Session $session
$scheduleList | format-table -autosize

# Associate the schedule with the data mover node
Set-DpHvBackupSchedule -Session $sess -ScheduleName "sched0" -Operation define `
    -DmNodesList hyperv1_HV_DM

# Remove the schedule association
Set-DpHvBackupSchedule -Session $sess -ScheduleName "sched0" -Operation remove `
    -DmNodesList hyperv1_HV_DM

Remove-DpHvSession -Session $sess

```

Example 12: Restore one or more VMs

Restore multiple VMs by referencing the backupIDs and restore them with new names and new restore destinations.

```

$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

# Restore a single VM with default parameters
$task = Restore-DpHvVm -Session $session -vmname "vm1"
$taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
while ("running" -eq $taskInfo.taskState) {
    start-sleep -seconds 30
    $taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
    if ($taskInfo.hasMoreData) {
        $results = Receive-DpHvTask -Session $session -TaskId $task.taskId
        write-verbose -verbose ("Started {0} Duration {1:g} Transferred `
            {2:N2} MB" -f $results.startTime, ((Get-Date)-$results.startTime),`
            ($results.totalBytesTransferred/1MB))
    }
}

$results = Receive-DpHvTask -Session $session -TaskId $task.taskId
$results

# restore multiple vms
$task = Restore-DpHvVm -Session $session -vmname vm1,vm2 -backupId 111111,222222 `
    -newVmName vm1_restored,vm2_restored -targetPath c:\restored,c:\restored
$taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
while ("running" -eq $taskInfo.taskState) {
    start-sleep -seconds 30
    $taskInfo = Get-DpHvTask -Session $session -TaskId $task.taskId
    if ($taskInfo.hasMoreData) {
        $results = Receive-DpHvTask -Session $session -TaskId $task.taskId
        write-verbose -verbose ("Started {0} Duration {1:g} Transferred `
            {2:N2} MB" -f $results.startTime, ((Get-Date)-$results.startTime),`
            ($results.totalBytesTransferred/1MB))
    }
}

$results = Receive-DpHvTask -Session $session -taskId $task.taskId
$results

# Get the restore history of VMs
$vmRestoreHistory = Get-DpHvVmRestoreTaskHistory -Session $session
$vmRestoreHistory

```



```
Remove-DpHvSession -Session $session
```

Example 13: Verify the configuration of Data Protection for Microsoft Hyper-V

After you run the configuration wizard, you can view the following configuration information by using the **Test-DpHvConfiguration** cmdlet:

- Information about the default data mover node such as the computer name, operating system, and location of the error log
- Information about the default mount proxy nodes such as the computer name, operating system, location of the error log, the state of the recovery agent, and the iSCSI status of the mount proxy nodes

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

$out1 = Test-DpHvConfiguration -session $session -nodetype DMNODE
$out2 = Test-DpHvConfiguration -session $session -nodetype MPNODE

Remove-DpHvSession -Session $session
```

Example 14: Verify the integrity of a VM backup by running cmdlets

After you back up a VM, you can verify the integrity of the VM backup by running instant access operations. You can automate the verification process by using the **InstantAccess** parameters in the **Restore-DpHvVm** cmdlet.

The following examples show the steps for verifying the integrity of a backup image for the VM named Ganymede.

1. Verify that the backup image of the VM named Ganymede can be used for restore operations. Run the following cmdlet to prepare a VM named Ganymede_verify for instant access operations and to automatically power on the new instant access VM.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

$vmName = "Ganymede"
$newVmName = "Ganymede_verify"
$restoreType = "InstantAccess"
$targetPath = "D:\SSD1\RecoveryAgent\mount"

Restore-DpHvVm -Session $session -VmName $vmName -NewVmName $newVmName `
  -TargetPath $targetPath -restoreType $restoreType -AutoStart

Remove-DpHvSession -Session $session
```

Important: If the VM that you are restoring is using a static IP address, a network IP address conflict will occur between the original VM and the temporary VM. To avoid this conflict, change the IP address of the temporary instant access VM after the VM is created.

2. To ensure that the VM Ganymede is running in instant access mode, run the **Get-DpHvIaVm** cmdlet to query the IBM Storage Protect server for instant access VMs.

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

Get-DpHvIaVm -VmName Ganymede -Session $session

Remove-DpHvSession -Session $session
```

3. Run a verification tool on the instant access VM named Ganymede_verify to verify that the backup image can be restored.

To verify the integrity of the disks and data, use a utility such as **chkdsk**, or a utility or application of your choice, to verify the virtual disks and data. If the temporary VM passes the integrity checks, you can remove the temporary resources that were created to support the instant access restore operation.

4. Remove the instant access VM named Ganymede_verify by running the following cmdlet:

```
$cred = Get-Credential -Message "Enter credentials" -UserName user_name
$session = New-DpHvSession -ComputerName computer_name -Credential $cred

$vmName = "Ganymede"
$newVmName = "Ganymede_verify"
$restoreType = "VmCleanup"
$targetPath = "D:\SSD1\RecoveryAgent\mount"

Restore-DpHvVm -Session $session -VmName $vmName -NewVmName $newVmName `
    -TargetPath $targetPath -restoreType $restoreType

Remove-DpHvSession -Session $session
```

Related reference

[“PowerShell cmdlets for Data Protection for Microsoft Hyper-V” on page 135](#)

Review the Data Protection for Microsoft Hyper-V cmdlets that you can use to protect your virtual machine (VM) data.

Chapter 9. Command reference

The following sections contain detailed information about the client commands that are used for IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V operations.

Issue these commands from the IBM Storage Protect backup-archive command line client. Start the command line client using either of the following methods on the Windows system:

- Go to **Start > Apps by name > IBM Storage Protect > Backup-Archive Command Line**.
- Open an Administrator command prompt window and change to the backup-archive client installation directory (**cd "C:\Program Files\tivoli\tsm\baclient"**). Run **dsmc.exe**.

To access the online help for the commands and options, issue the following command:

```
dsmc help
```

The online help contains information that applies to the IBM Storage Protect backup-archive client, IBM Storage Protect for Virtual Environments: Data Protection for VMware, and Data Protection for Microsoft Hyper-V.

Reading syntax diagrams

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.

- The **▶▶—** symbol indicates the beginning of a syntax diagram.
- The **—▶** symbol at the end of a line indicates that the syntax diagram continues on the next line.
- The **▶—** symbol at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The **—▶◀** symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or a variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element)

Symbols

Enter these symbols *exactly* as they appear 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 *<var_name>* indicate variables. In this example, you can specify a *<var_name>* when you enter the **cmd_name** command.


➤ cmd_name — *<var_name>* ➤

Repetition

An arrow returning to the left means that the item can be repeated. A character within the arrow means that you must separate repeated items with that character.



A footnote (1) by the arrow refers to a limit that tells how many times the item can be repeated.




Notes:

¹ Specify *repeat* up to 5 times.

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 this example, you must choose A, B, or C.




Optional choices

When an item is *below* the line, that item is optional. In the first example, you can select A or nothing at all.




When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.



Repeatable choices

A stack of items followed by an arrow returning to the left indicates that you can select more than one item, or in some cases, repeat a single item.


In this example, you can select any combination of A, B, or C.



Defaults

Defaults are above the line. The default is selected unless you override it, or you can select the default explicitly. To override the default, include an option from the stack below the line.

In this example, A is the default. Select either B or C to override A.

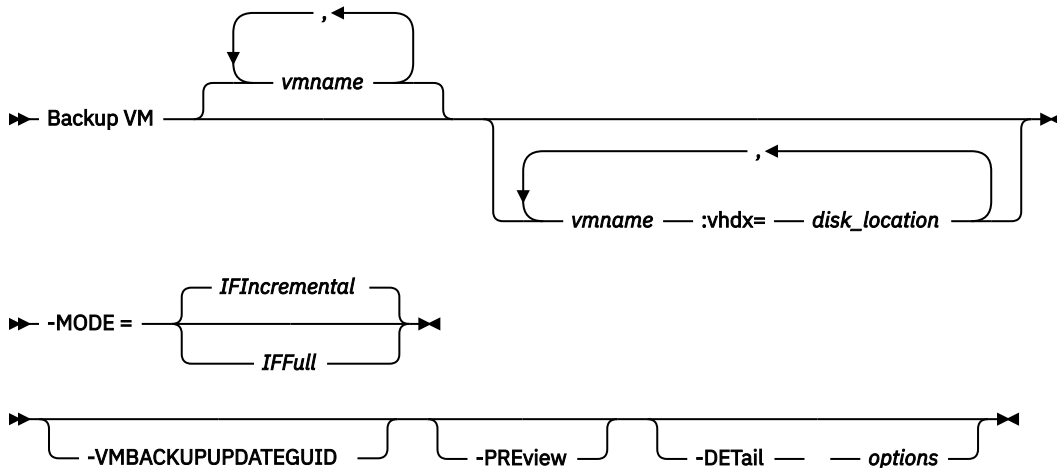


Backup VM

Use the **backup vm** command to back up Hyper-V virtual machines.

You can back up Hyper-V guests that exist on a local disk, a SAN-attached disk, a Cluster Shared Volume (CSV), or guests that exist on a remote file server share. Remote file server shares must be on a Windows Server 2012 or later system. In addition, remote file shares must be Server Message Block (SMB) 3.0 with the File Server VSS Agent Service installed on the server.

Syntax



Parameters

vmname

Specifies the name of the virtual machine that you want to back up. To specify multiple virtual machines, separate multiple virtual machine names with commas (vm1, VM2, Vm5), or use the domain.vmfull option. The names are case-sensitive and must match the capitalization that is shown on the Hyper-V host in the **Hyper-V Manager > Virtual Machines** view.

Wildcards can be used in virtual machine names.

Restriction: Do not use the word "aggregate" as a VM name. The word "aggregate" is reserved for the VM name field in an IBM Storage Protect message. Messages that contain the word "aggregate" as the VM name show statistics that represent the total value of all VM guests that are backed up per data mover.

vmname:vhdxdisk_location

This parameter specifies the virtual machine hard disk (VHDX) to include in Hyper-V virtual machine backup operations.

The *vmname* variable specifies the name of the VM to back up. Wildcard characters can be used to select VM names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character.

The **:vhdxdisk_location** keyword specifies the location of the VM disk to include in the backup operation. The disk location is specified in the format "*controller_type controller_number disk_location_number_inside_controller*". The controller type must be "SCSI" or "IDE". For example:

```
dsmc backup vm "vm1:VHDX=IDE 1 0"
```

You can obtain the disk location information in the **Hyper-V Manager**. In the **Virtual Machines** view, right-click a VM and click **Settings**. In the **Hardware** section of the **Settings** window, locate the IDE Controller or SCSI Controller, and click **Hard Drive** to view the hard disk settings. The controller number and disk location are displayed in the **Controller** and **Location** fields. You can also obtain the disk location information by running the Windows PowerShell cmdlet **Get-VMHardDiskDrive**.

Restriction: Do not use the word "aggregate" as a VM name. The word "aggregate" is reserved for the VM name field in an IBM Storage Protect message. Messages that contain the word "aggregate" as the VM name show statistics that represent the total value of all VM guests that are backed up per data mover.

You can exclude a VM disk from backup operations by specifying the exclude operator (-) before the **vhdxdisk_location** keyword. For example, use **-vhdxdisk_location** to exclude a VM disk from the backup operation of a VM:

```
dsmc backup vm "vm1:-VHDX=IDE 1 0:-VHDX=SCSI 0 1"
```

If you specify multiple VM disks to include or exclude, the **vhdxdisk_location** or **-vhdxdisk_location** keyword and associated values must be separated by colons, with no intervening space characters. For example:

```
dsmc backup vm "*:-VHDX=IDE 1 0:-VHDX=SCSI 0 1"
```

If you specify multiple VM names and VM disks, the VM name and associated values must be separated by semicolons, with no intervening space characters. For example:

```
dsmc backup vm "vm1:-VHDX=IDE 1 0:-VHDX=SCSI 0 1;vm2:-VHDX=IDE 1 0:-VHDX=SCSI 0 1"
```

-MODE

You must specify the backup mode to use when backing up a virtual machine by adding the **-mode** parameter on the command line. The following modes can be specified:

IFFull

Incremental-forever-full mode. In this mode, a snapshot of all used blocks on a virtual machine's disks are backed up to the server. The backup includes configuration information, and all of the disks.

IFIncremental

Incremental-forever-incremental. In this mode, a snapshot is created of the blocks that have changed since the last incremental forever backup operation, whether full or incremental. The backup includes configuration information, and all of the disks. This value is the default.

-VMBACKUPUPDATEGUID

This option updates the globally unique identifier (GUID) for the virtual machine that you are backing up. This parameter is intended for use only in the following scenario:

You want to restore an already backed up virtual machine named ORION. But, before you shut down and replace the copy of ORION that is running in your production environment, you want to verify the configuration of the restored virtual machine before you use it to replace the existing ORION.

1. You restore the ORION virtual machine and give it a new name: `dsmc restore vm Orion -vmname=Orion2`
2. You update and verify the ORION2 virtual machine and determine that it is ready to replace the existing virtual machine that is named ORION.
3. You power down and delete ORION.
4. You rename ORION2 so it is now named ORION.
5. The next time that you back up ORION, by using either an incremental-forever full, or incremental-forever-incremental backup, you add the **-VMBACKUPUPDATEGUID** parameter to the **backup vm** command. This option updates the GUID, on the IBM Storage Protect server, so the new GUID is associated with the stored backups for the ORION virtual machine. The chain of incremental backups is preserved; there is no need to delete existing backups and replace them with new backups.

-PREview

This parameter displays additional information about a virtual machine, including the labels of the virtual hard disks that are in the virtual machine.

When you issue the `-preview` option, the backup operation does not start. You must issue the backup command without the `-preview` option to start the backup operation.

You can use both the `-preview` option and the `-detail` option on the command to display information about subdisks that are included when the backup is run. A subdisk is the AVHDX file that is created when a snapshot is taken of a VHDX file.

-DEtail

This parameter displays detailed information about a virtual machine. Use this option with `-preview` to view more details about the disks that are involved in the backup operation.

When you issue the `-detail` option, the backup operation does not start. You must issue the backup command without the `-detail` option to start the backup operation.

Return codes for virtual machine backup operations

Backup operations for virtual machines can complete with the return codes that are shown in the following table.

<i>Table 15. Return codes from backup vm commands</i>	
Return code	Description
0	A command to back up one or more virtual machines completed successfully.
8	A command to back up multiple virtual machines succeeded for only some of the virtual machines that were targeted by the command. Examine the log file to determine the processing status for each of the targeted virtual machines.
12	Indicates that either of the following error conditions occurred: <ul style="list-style-type: none"> • The backup command could not back up any of the virtual machines that were targets of the backup operation. • The backup command failed and it stopped before all virtual machines that were specified were inspected. Examine the log file to determine the reason for the failure.

Tip about the final statistics: If you are running multiple backup sessions, the value that is displayed in the **Data transfer time** field in the final statistics can be higher than the value in the **Elapsed processing time** field. The data transfer time is the sum of the times that each backup takes to send data across

the network. This number does not include the time for the data mover to read the data from disk before sending it, nor the time to wait for server transactions to complete. This number can be greater than the elapsed processing time if the operation uses multiple concurrent sessions to move data, such as multi-session backup operations. This value includes the time that it takes to send data more than once due to retries, such as when a file changes during a backup operation.

Example commands

1. The following command starts an incremental-forever incremental backup of a Hyper-V virtual machine that is named "VM1":

```
dsmc backup vm VM1 -mode=ifincremental
```

2. For Windows Server 2016 or later operating systems: The following command excludes an IDE disk (with controller number 1 and disk location 0) and a SCSI disk (with controller number 0 and disk location 1) from an incremental-forever incremental RCT backup of a virtual machine, "vm2":

```
dsmc backup vm "vm2:-VHDX=IDE 1 0:-VHDX=SCSI 0 1" -mode=ifincremental
```

3. For Windows Server 2016 or later operating systems: The following command shows the preview of an incremental-forever incremental RCT backup of a virtual machine, "VM05":

```
dsmc backup vm VM05 -mode=ifincremental -preview
```

In the command output, the -preview parameter displays the VHDX labels in the virtual machine. When the -detail parameter is specified with the -preview parameter, no additional information is shown for Hyper-V RCT backups.

```
Backup VM command started. Total number of virtual machines to process: 1
```

1. VM Name: VM05

```
Domain Keyword:      VM05
Mode:                Incremental Forever - Incremental
Target Node Name:    NODE14
Data Mover Node Name: NODE14
Cluster Resource:    No
```

```
Disk[1]
Name: \\node14\d$\Hyper_V_Virtual_Machine\VM05\Virtual Hard Disks\VM05.vhdx
Capacity:            15.00 GB
Size:                10.91 GB
Status:              included
Disk Type:           VHDX
Number of Subdisk:   0
Controller Location: IDE 0 0
```

```
Disk[2]
Name: \\node14\d$\Hyper_V_Virtual_Machine\VM05\Virtual Hard Disks\
      VM05_Disk2.vhdx
Capacity:            2.00 GB
Size:                132.00 MB
Status:              included
Disk Type:           VHDX
Number of Subdisk:   0
Controller Location: SCSI 0 1
```

```
Total number of virtual machines processed: 1
```

4. For Windows Server 2012 or 2012 R2: The following command starts an incremental forever-incremental backup of a Hyper-V virtual machine, "VM03":

```
dsmc backup vm VM03 -mode=ifincremental -preview
```

In the command output, the -preview parameter displays the VHDX labels in the virtual machine:

1. VM Name: VM03

```
Domain Keyword:      all-vm
Mode:                Incremental Forever - Incremental
Target Node Name:    NODE14_HV_DM
Data Mover Node Name: NODE14_HV_DM
Cluster Resource:    No

Disk[1]
Name: \\NODE14\d$\Hyper-V\VM03\VM03\Virtual Hard Disks\VM03.vhdx
Capacity:           64.00 GB
Size:               28.91 GB
Status:             excluded
Disk Type:          VHDX
Number of Subdisk:  1
```

When the `-detail` parameter is specified with the `-preview` parameter, the VHDX labels and their subdisks are shown. The following example output is abbreviated to show only information about one virtual machine and one disk:

1. VM Name: VM03

```
Domain Keyword:      all-vm
Mode:                Incremental Forever - Incremental
Target Node Name:    NODE14_HV_DM
Data Mover Node Name: NODE14_HV_DM
Cluster Resource:    No

Disk[1]
Name: \\NODE14\d$\Hyper-V\VM03\VM03\Virtual Hard Disks\VM03.vhdx
Capacity:           64.00 GB
Size:               28.91 GB
Status:             excluded
Disk Type:          VHDX
Number of Subdisk:  1

Subdisk[1]
Name: \\NODE14\d$\Hyper-V\VM03\VM03\Virtual Hard Disks\
      VM03_94F6257B-5C61-45F1-BD62-3323DCF26954.avhdx
Capacity:           64.00 GB
Size:               180.00 MB
Status:             excluded
Disk Type:          AVHDX
```

Options file examples

The `domain.vmfull` option is used to process specific virtual machines. In the following example, the `domain.vmfull` option is specified as follows:

```
domain.vmfull VM04,VM05
```

The following command shows a preview of a full backup of virtual machines specified by the `domain.vmfull` option. The command displays preview information about each virtual machine:

```
dsmc backup vm -mode=iffull -preview
```

The following output is shown on Windows Server 2016 and later operating systems:

```
Backup VM command started. Total number of virtual machines to process: 2
```

1. VM Name: VM04

```
Domain Keyword:      VM04
Mode:                Incremental Forever - Full
Target Node Name:    NODE14
Data Mover Node Name: NODE14
Cluster Resource:    No

Disk[1]
Name: \\node14\d$\Hyper_V_Virtual_Machine\VM04\Virtual Hard Disks\VM04.vhdx
Capacity:           36.00 GB
Size:               9.16 GB
Status:             included
```

```

Disk Type:          VHDX
Number of Subdisk:  0
Controller Location: IDE 0 0

2. VM Name: VM05

Domain Keyword:     VM05
Mode:               Incremental Forever - Full
Target Node Name:   NODE14
Data Mover Node Name: NODE14
Cluster Resource:   No

Disk[1]
Name: \\node14\d$\Hyper_V_Virtual_Machine\VM05\Virtual Hard Disks\VM05.vhdx
Capacity:          15.00 GB
Size:              10.91 GB
Status:            included
Disk Type:         VHDX
Number of Subdisk: 0
Controller Location: IDE 0 0

Disk[2]
Name: \\node14\d$\Hyper_V_Virtual_Machine\VM05\Virtual Hard Disks\
VM05_Disk2.vhdx
Capacity:          2.00 GB
Size:              132.00 MB
Status:            included
Disk Type:         VHDX
Number of Subdisk: 0
Controller Location: SCSI 0 1

```

Total number of virtual machines processed: 2

The following output is shown on Windows Server 2012 and 2012 R2:

Backup VM command started. Total number of virtual machines to process: 2

```

1. VM Name: VM04

Domain Keyword:     all-vm
Mode:               Incremental Forever - Incremental
Target Node Name:   NODE14_HV_DM
Data Mover Node Name: NODE14_HV_DM
Cluster Resource:   No

Disk[1]
Name: \\NODE14\d$\Hyper-V\VM04\VM04\Virtual Hard Disks\VM04.vhdx
Capacity:          64.00 GB
Size:              28.91 GB
Status:            excluded
Disk Type:         VHDX
Number of Subdisk: 1

Subdisk[1]
Name: \\NODE14\d$\Hyper-V\VM04\VM04\Virtual Hard Disks\
VM04_94F6257B-5C61-45F1-BD62-3323DCF26954.avhdx
Capacity:          64.00 GB
Size:              180.00 MB
Status:            excluded
Disk Type:         AVHDX

2. VM Name: VM05

Domain Keyword:     all-vm
Mode:               Incremental Forever - Incremental
Target Node Name:   NODE14_HV_DM
Data Mover Node Name: NODE14_HV_DM
Cluster Resource:   No

Disk[1]
Name: \\NODE14\d$\Hyper-V\disks\Windows 10.vhdx
Capacity:          20.00 GB
Size:              18.75 GB
Status:            excluded
Disk Type:         VHDX
Number of Subdisk: 1

Subdisk[1]
Name: \\NODE14\d$\Hyper-V\disks\
Windows 10_15F8A5AA-C104-4C74-8F68-B57B27592F8A.avhdx

```

```

Capacity:          20.00 GB
Size:              112.00 MB
Status:            excluded
Disk Type:         AVHDX

Disk[2]
Name: \\NODE14\\e$\\Hyper-V\\disks\\Windows10_disk2\\Windows10_disk2.vhdx
Capacity:          5.00 GB
Size:              5.00 GB
Status:            excluded
Disk Type:         VHDX
Number of Subdisk: 1

Subdisk[1]
Name: \\NODE14\\e$\\Hyper-V\\disks\\Windows10_disk2\\
      Windows10_disk2_15F8A5AA-C104-4C74-8F68-B57B27592F8A.avhdx
Capacity:          5.00 GB
Size:              4.00 MB
Status:            excluded
Disk Type:         AVHDX

Disk[3]
Name: \\NODE14\\e$\\Hyper-V\\disks\\Windows10_disk2\\Windows10_disk5.vhdx
Capacity:          1.00 GB
Size:              1.00 GB
Status:            included
Disk Type:         VHDX
Number of Subdisk: 1

Subdisk[1]
Name: \\NODE14\\e$\\Hyper-V\\disks\\Windows10_disk2\\
      Windows10_disk5_15F8A5AA-C104-4C74-8F68-B57B27592F8A.avhdx
Capacity:          1.00 GB
Size:              4.00 MB
Status:            included
Disk Type:         AVHDX

Total number of virtual machines processed: 2
ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

```

Related reference

[“Detail” on page 169](#)

Use the `detail` option to display management class, file space, and backup information.

[“Domain.vmfull” on page 169](#)

The `domain.vmfull` option specifies the virtual machines (VMs) to include in your full virtual machine image backup operations.

[“Mbobjrefreshthresh” on page 183](#)

The `mbobjrefreshthresh` (megablock object refresh threshold) option is a number defining a threshold. When the number of IBM Storage Protect objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

[“Mbpctrefreshthresh” on page 183](#)

The `mbpctrefreshthresh` (megablock percentage refresh threshold) option is a number defining a threshold. When the number of IBM Storage Protect percentage of objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

[“Mode” on page 182](#)

Use the `mode` option to specify the backup mode to use when performing specific backup operations.

[“Query VM” on page 154](#)

Use the **query VM** command to list and verify the successful backups of virtual machines (VMs).

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

Expire

Use the **expire** command to deactivate the current backup of a virtual machine (VM) on the IBM Storage Protect server.

When you are working in interactive mode, a prompt notifies you before objects are expired.

The **expire** command does not remove the VM from the local host. If you expire a VM that still exists on your local host, the VM is backed up again during the next incremental backup, unless you exclude the VM from backup processing.

Syntax

➔ EXPire — — OBJTYPE=VM — — vmname — — options ➔

Parameters

OBJTYPE=VM vmname

vmname specifies the name of a VM. The active backup for the specified VM is expired on the IBM Storage Protect server. The VM name cannot contain wildcard characters.

When objtype=VM is specified, the expire command expires only full VM backups (MODE=IFFULL) for the VM that is specified on the vmname parameter.

Table 16. Expire command: Related options

Option	Where to use
dateformat “Dateformat” on page 167	Client options file (dsm.opt) or command line.
noprompt “Noprompt” on page 184	Command line only.
numberformat “Numberformat” on page 184	Client options file (dsm.opt) or command line.
pick “Pick” on page 186	Command line only.
timeformat “Timeformat” on page 188	Client user-options file (dsm.opt) or command line.

Example

Task

Deactivate the current backup of the VM that is named vm_test.

Command: dsmc expire -objtype=VM vm_test

Query VM

Use the **query VM** command to list and verify the successful backups of virtual machines (VMs).

Query VM for Microsoft Hyper-V virtual machines

Use the **query vm** command to determine which Hyper-V virtual machines were backed up.

Supported Clients

This command is valid on Windows clients that are installed on a Hyper-V host system.

Syntax



Parameters

vmname

Specifies the virtual machine host name that you want to query. The virtual machine name is case-sensitive. If you specify a VM name on the command, the name cannot contain wildcard characters.

If you omit the virtual machine name, the command displays all VM backups on the IBM Storage Protect server.

Table 17. Query VM command: Related options for Hyper-V virtual machine queries.

Option	Where to use
detail	Command line. Displays the details of each disk (label, name) and its status (protected or excluded), and incremental-forever backup performance statistics.
inactive	Command line.
pitdate	Command line.
pittime	Command line.
vmrestoretype	Command line.

Examples

Task

List all virtual machines that have been backed up by Data Protection for Microsoft Hyper-V on the Hyper-V host.

```
dsmc query vm
```

Query VM examples (Hyper-V)

The following example shows a **query VM** command that displays summary information about all Hyper-V virtual machines that have been backed up.

```
dsmc query vm
```

Query Virtual Machine for Full VM backup

# Machine	Backup Date	Mgmt Class	Size	Type	A/I Location	Virtual
1	03/19/2017 17:54:34	STANDARD	17.00 GB	IFFULL	A SERVER	DeptA_VM05
2	03/20/2017 01:51:34	STANDARD	15.00 GB	IFINCR	A SERVER	
3	03/20/2017 01:46:19	STANDARD	36.00 GB	IFFULL	A SERVER	DeptA_VM04

The following **query VM** command with the **-detail** option shows detailed information about Hyper-V VMs that have been backed up. The detailed output includes the type of backup that was performed, the size of the virtual machine, information about its disks, and statistics.

```
dsmc query vm -detail
```

Query Virtual Machine for Full VM backup

#	Backup Date	Mgmt Class	Size	Type	A/I	Location	Virtual Machine
1	03/19/2017 17:54:34	STANDARD	17.00 GB	IFFULL	A	SERVER	DeptA_VM05
The size of this incremental backup: n/a The number of incremental backups since last full: 0 The amount of extra data: 0 The IBM Storage Protect objects fragmentation: 0 Backup is represented by: 99 IBM Storage Protect objects Application protection type: n/a Backup is compressed: No Backup is deduplicated: No Snapshot type: Hyper-V RCT Application Consistent Disk[1]Name: DeptA_VM05.vhdx Disk[1]Location: IDE 0 0 Disk[1]Status: Protected Disk[2]Name: DeptA_VM05_Disk2.vhdx Disk[2]Location: SCSI 0 1 Disk[2]Status: Protected Disk[3]Name: Disk 7 2.00 GB Bus 0 Lun 4 Target 0 Disk[3]Location: SCSI 0 0 Disk[3]Status: Skipped: Physical disk Disk[4]Name: Disk 8 2.50 GB Bus 0 Lun 5 Target 0 Disk[4]Location: SCSI 0 2 Disk[4]Status: Skipped: Physical disk							
2	03/20/2017 01:51:34	STANDARD	15.00 GB	IFINCR	A	SERVER	DeptA_VM_W2k08R2
The size of this incremental backup: 544.00 KB The number of incremental backups since last full: 1 The amount of extra data: 0 The IBM Storage Protect objects fragmentation: 2 Backup is represented by: 37 IBM Storage Protect objects Application protection type: n/a Backup is compressed: No Backup is deduplicated: No Snapshot type: Hyper-V RCT Crash Consistent Disk[1]Name: DeptA_VM_W2k08R2.vhdx Disk[1]Location: IDE 0 0 Disk[1]Status: Protected							
3	03/20/2017 01:46:19	STANDARD	36.00 GB	IFFULL	A	SERVER	DeptA_VM04
The size of this incremental backup: n/a The number of incremental backups since last full: 0 The amount of extra data: 0 The IBM Storage Protect objects fragmentation: 0 Backup is represented by: 79 IBM Storage Protect objects Application protection type: n/a Backup is compressed: No Backup is deduplicated: No Snapshot type: Hyper-V RCT Application Consistent Disk[1]Name: DeptA_VM04.vhdx Disk[1]Location: IDE 0 0 Disk[1]Status: Protected							
----- All averages are calculated only for incremental forever backups displayed above. The average size of incremental backup: 544.00 KB The average number of incremental backups since last full: 0 The average overhead of extra data: 0 The average objects fragmentation: 0 The average number of objects per backup: 71							

The detailed output also includes the snapshot type and disk information such as the following information:

Snapshot type

The type of snapshot that was taken during the VM backup operation:

Hyper-V RCT Application Consistent

A quiesced snapshot that was created with Hyper-V Resilient change Tracking (RCT) on Windows Server 2016.

Hyper-V RCT Crash Consistent

A non-quieted snapshot that was created with Hyper-V RCT on Windows Server 2016.

Hyper-V VSS

A snapshot that was created with Volume Shadow Copy Service (VSS) on Windows Server 2012 or Windows Server 2012 R2.

Disk[n]Location

The disk location of VM disk *n*, where *n* is a number. The disk location consists of the disk controller type, "IDE" or "SCSI", followed by the controller number and device location number.

Disk[n]Status

The backup status of VM disk *n*, where *n* is a number.

Protected

Indicates that the data on the VM disk is backed up.

Skipped: Excluded by user

Indicates that the VM disk is excluded during backup operations as specified by the `exclude.vmdisk` option.

Skipped: Physical disk

Indicates that the VM disk is a physical disk (pass-through disk) and its data is not backed up. Only the disk configuration information is backed up.

The following example shows the syntax to use to list detailed output for a specific virtual machine named DeptA_VM_W2k08R2.

```
dsmc query vm DeptA_VM_W2k08R2 -detail
```

Query Virtual Machine for Full VM backup

#	Backup Date	Mgmt Class	Size	Type	A/I	Location	Virtual Machine
1	03/20/2017 01:51:34	STANDARD	15.00 GB	IFINCR	A	SERVER	DeptA_VM_W2k08R2

The size of this incremental backup: 544.00 KB
The number of incremental backups since last full: 1
The amount of extra data: 0
The IBM Storage Protect objects fragmentation: 2
Backup is represented by: 37 IBM Storage Protect objects
Application protection type: n/a
Backup is compressed: No
Backup is deduplicated: No
Snapshot type: Hyper-V RCT Crash Consistent
Disk[1]Name: Jimmy_VM_Windows2008R2.vhdx
Disk[1]Location: IDE 0 0
Disk[1]Status: Protected

All averages are calculated only for incremental forever backups displayed above.

The average size of incremental backup: 544.00 KB
The average number of incremental backups since last full: 1
The average overhead of extra data: 0
The average objects fragmentation: 2
The average number of objects per backup: 37

The following command displays all temporary VMs that are running in instant access mode that were created by an instant access operation.

```
q vm * -vmrestoretype=instantaccess
```

Related reference

[“Exclude.vmdisk” on page 172](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

[“Vmprocessvmwithphysdisks” on page 206](#)

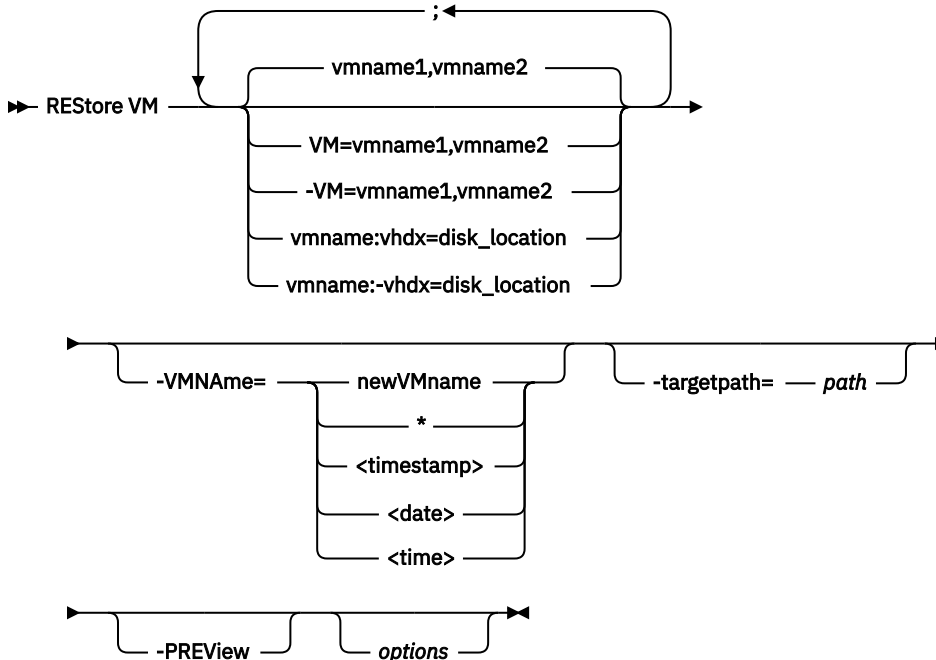
Use the `vmprocessvmwithphysdisks` option to control whether Hyper-V RCT virtual machine (VM) backups are processed if the VM has one or more physical disks (pass-through disks) provisioned.

Restore VM

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

If the VM that you are restoring exists on the Hyper-V host server, it is shut down and deleted before it is restored from the image that is stored on the IBM Storage Protect server. The restore operation then creates the VM such that its content and configuration is identical to what it was when the backup occurred. Even though the VM is shut down before it is deleted, manually shutting down the VM before you run the **restore vm** command is a good practice to bring any in-progress application activities to an orderly stop.

Syntax



Parameters

Any parameter that contains spaces must be enclosed in quotation marks (" ").

vmname

Specify the name of one or more VMs that you want to restore. The name is the VM display name. Separate multiple VM names with commas (for example, `vm1, vm2, vm5`). If you backed up template VMs, the `vmname` parameter can specify the name of a template VM to restore.

Wildcard characters can be used to select VMs names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character. For example:

- `restore vm VM_TEST*` restores all VMs that have names that begin with "VM_TEST".
- `restore vm VM??` restores any VM that has a name that begins with the letters "VM", followed by 2 characters.

You must specify one or more VMs to restore.

vm=vmname

The `vm=` keyword specifies that the next set of values is a list of VM names. The `vm=` keyword is the default and is not required.

Wildcard characters can be used in VM names. For the specification of the *vmname* parameter, see [“vmname” on page 158](#).

In the following example, `vm=` is specified and commas are used to separate two machine names.

```
restore vm vm=my_vm1,my_vm2
```

-vm=vmname

You can exclude a VM from a restore operation by specifying the exclude operator (-) before the `vm=` keyword.

Use the `-vm=` keyword to exclude a list of VMs from a larger group of VM backups, such as a group of VMs that begin with a VM name pattern. For example, if you want to restore all VMs that start with `Dept99_` but prevent `vm2` from being restored, issue the following command:

```
restore vm vm=Dept99_*;-vm=Dept99_vm2
```

Wildcard characters can be used with the `-vm=` keyword to exclude VM names that match a pattern, for example:

- Exclude all files that have `test` in the host name:

```
-vm=*test*
```

- Include all VMs with names such as: `test20`, `test25`, `test29`, and `test2A`:

```
vm=test2?
```

Restriction: You cannot use the exclude operator (-) to exclude a VM host domain. The exclude operator works only at the VM name level.

vmname:vhdxdisk_location

This parameter specifies the VM hard disk (VHDX) to include in Hyper-V VM restore operations. Specify this option only if you want to restore one or more specific disks, but not all disks.

The *vmname* variable specifies the name of the VM to restore. Wildcard characters can be used to select VM names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character.

The **:vhdxdisk_location** keyword specifies the location of the VM disk to include in the restore operation. The disk location is specified in the format *"controller_type controller_number disk_location_number_inside_controller"*. The controller type must be "SCSI" or "IDE". For example:

```
dsmc restore vm "vm1:vhdxdisk_location=IDE 1 0"
```

If you specify multiple VM disks to include, the **vhdxdisk_location** keyword and associated values must be separated by colons, with no intervening space characters. For example:

```
dsmc restore vm "*:vhdxdisk_location=IDE 1 0:vhdxdisk_location=SCSI 0 1"
```

If you specify multiple VM names and VM disks, the VM name and associated values must be separated by semicolons, with no intervening space characters. For example:

```
dsmc restore vm "vm1:vhdxdisk_location=IDE 1 0:vhdxdisk_location=SCSI 0 1;vm2:vhdxdisk_location=IDE 1 0:vhdxdisk_location=SCSI 0 1"
```

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 use the **-preview** parameter to identify the original disk name and capacity.

- The existing disk must be at least as large as the disk that you want to restore.
- The existing disk location must be the same as the disk that 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.

vmname:-vhdx=disk_location

The *vmname* variable specifies the name of the VM to restore. Wildcard characters can be used to select VM names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character.

This **-vhdx=disk_location** option is used to specify the disk location of one or more virtual disks to exclude from restore operations. For example, use **-vhdx=** to exclude a VM disk from the restore operation of a VM:

```
dsmc restore vm "vm1:-vhdx=IDE 1 0:-vhdx=SCSI 0 1"
```

If you specify multiple VM disks to exclude, the **-vhdx=** keyword and associated values must be separated by colons, with no intervening space characters. For example:

```
dsmc restore vm "*:-vhdx=IDE 1 0:-vhdx=SCSI 0 1"
```

If you specify multiple VM names and VM disks, the VM name and associated values must be separated by semicolons, with no intervening space characters. For example:

```
dsmc restore vm "vm1:-vhdx=IDE 1 0:-vhdx=SCSI 0 1;vm2:-vhdx=IDE 1 0:-vhdx=SCSI 0 1"
```

-VMName=

Specifies the new name for the VM after it is restored, if you do not want to use the name specified by the **VM=** parameter. Ensure that you enclose the new VM name in double quotation marks.

newVMname

Specify a new VM name to use for the restored VM.

The following characters are not supported in names of restored VMs:

```
: ; ' \ / " ? , < > |
```

A restore command that includes unsupported characters will fail with error message ANS9117E.

The * (asterisk) symbol is a metacharacter that represents the original name of the VM that is being restored. Placing valid characters before or after the asterisk creates a prefix or suffix in the name of the restored VM.

The following characters are not supported in names of restored VMs:

```
: ; ' \ / " ? , < > |
```

A restore command that includes unsupported characters will fail with error message ANS9117E.

You can use the * symbol in the following manner:

- Use the original VM name for the restored VM name by specifying **vmname="*"**.
- Append a suffix to the original VM name for the restored VM. For example, if the original VM names are VM1 and VM2, you can append the suffix "_restored" to the original VM names by specifying the following command:

```
dsmc restore vm VM1,VM2 -VMName="*_restored"
```

The names of the restored VMs are VM1_restored and VM2_restored.

- Insert a prefix before the original VM name for the restored VM. For example, if the original VM name is VM3, you can insert the prefix "new_" to VM3 by specifying the following command:

```
dsmc restore vm VM3 -vmname="new_★"
```

The name of the restored VM is new_VM3.

<timestamp>

Appends a time stamp with the date and time of the restore operation to the name of the restored VM. The <timestamp> parameter is a keyword, and must include the bracket symbols ("<" and ">").

The format for the time stamp string is determined by the DATEFORMAT and TIMEFORMAT options in the dsm.opt file. A dash is used as a delimiter for the time stamp that is returned by the <timestamp> parameter.

For example, to restore two VMs named VM5 and VM6, and append the date and time of restore to the restored VM names, issue the following command:

```
dsmc restore vm VM5,VM6 -vmn="★_<timestamp>"
```

The names of the restored VMs are VM5_06-22-2017_14-56-55 and VM6_06-22-2017_14-56-55.

<date>

Appends the date of the restore operation to the name of the restored VM. The <date> parameter is a keyword, and must include the bracket symbols ("<" and ">").

The format of the date string is determined by the DATEFORMAT option in the dsm.opt file. A dash is used as a delimiter for the date that is returned by the <date> parameter.

For example, to insert the prefix "new_" before the VM named VM3, and append the restore date to the restored VM name, issue the following command:

```
dsmc restore vm VM3 -vmname="new_★_<date>"
```

The name of the restored VM is new_VM3_06-22-2017.

<time>

Appends the time of the restore operation to the name of the restored VM. The <time> parameter is a keyword, and must include the bracket symbols ("<" and ">").

The format of the time string is determined by the TIMEFORMAT option in the dsm.opt file. A dash is used as a delimiter for the time that is returned by the <time> parameter.

For example, to append the suffix "_today_" after the VM named VM8, and add the restore time to the restored VM name, issue the following command:

```
dsmc restore vm VM8 -vmn="★_today_<time>"
```

The name of the restored VM is VM8_today_14-56-55.

-targetpath=

Specifies the path that you want to restore the VM to.

This parameter is required if the -vmname parameter is used and optional otherwise. Use this parameter to restore the VM to an alternative path.

-PREVIEW

Use this parameter to verify the results of a restore operation without restoring VMs. The **-preview** parameter provides a list of VMs that will be restored and information about the VMs, such as the names of virtual hard disks and the size of the disks.

When you issue the **-preview** parameter with the **restore vm** command, the restore operation does not start. The restore operation starts only if the **-preview** parameter is removed from the command.

For more information, see [“Preview virtual machine restore operations” on page 164](#).

Table 18. Restore VM command: Related options when restoring Hyper-V virtual machines

Option	Where to use
inactive	Command line
pick	Command line
pitdate	Command line
pittime	Command line
replace	Command line, client options file, or client preferences editor.
vmbackdir	Command line, client options file.
vmmaxrestoreparalleldisks	Command line, options file
vmmaxrestoresessions	Command line, options file
vmmaxrestoreparallelvm	Command line, options file
vmrestoretype	Command line
vmautostartvm	Command line or client options file.

Tip: For a host that is in a cluster, if you restored a deleted VM or if you restored a VM with a new VM name, you must configure the restored VM for high availability by using Microsoft Failover Cluster Manager, System Center Virtual Machine Manager, or PowerShell cmdlets. For instructions about configuring a VM for high availability, see the Microsoft documentation.

Tip about the final statistics: If you are running multiple restore sessions, the value that is displayed in the **Data transfer time** field in the final statistics can be higher than the value in the **Elapsed processing time** field. The data transfer time is the sum of the times that each restore operation takes to send data across the network. This number does not include the time for the data mover to read the data from disk before sending it, nor the time to wait for server transactions to complete. This number can be greater than the elapsed processing time if the operation uses multiple concurrent sessions to move data, such as multi-session restore operations. This value includes the time that it takes to send data more than once due to retry operations.

Examples

Task

Restore the most recent backup version of a virtual machine named VM1 to the drive and path it was in when it was backed up.

```
dsmc restore vm VM1 -vmrestoretype=noninstant
```

Task

Restore the most recent backup version of a virtual machine named vm1 to the drive and path it was in when it was backed up. Replace the existing virtual machine without prompting.

```
dsmc restore vm vm1 -replace=yes
```

Task

Restore the backed up virtual machine named VM1 to a new name (vm2):

```
dsmc restore vm VM1 -VmName=vm2
```

Task

Restore the backed up virtual machine named vm1, and assign it a new name (vm2). Issue a prompt before overwriting vm2, if that virtual machine already exists.

```
dsmc restore vm vm1 -VmName=vm2 -replace=prompt
```

Task

Restore the virtual machine named vm1 to a specific drive and path, without renaming the virtual machine:

```
dsmc restore vm vm1 -targetpath="E:\New Path"
```

Task

Restore the virtual machine named vm1, but rename it as vm2 and restore it to a new path:

```
dsmc restore vm vm1 -VmName=vm2 -targetpath=F:\NewPath
```

Task

Use `-pick` and `-inactive` to show active and inactive backups for a virtual machine that is named vm1. You choose the backup to restore from a list:

```
dsmc restore vm vm1 -pick -inactive
```

Task

Run an instant access operation from the command line. For instructions, see [Chapter 7, “Running instant access operations,”](#) on page 125.

Task

From a list of VMs, restore four VMs in parallel to the `c:\hypervrestore` folder. In the restore operation, allocate five sessions per VM that is being restored, with a maximum of two disks restored in parallel per VM. Restore the VMs to VM names that are appended with the time and date of the restore operation.

```
dsmc restore vm vm1,vm2,vm3,vm4,vm5,vm6,vm7,vm8 -vmmaxrestoreparallelvm=4  
-vmmaxrestoresessions=20 -vmmaxrestoreparalleldisks=2  
-vmname="*_<timestamp>" -targetpath=c:\hypervrestore
```

Task

Of all the VMs that match the pattern `vm*`, restore four VMs in parallel to the `c:\hypervrestore` folder. In the restore operation, allocate five sessions per VM that is being restored, with a maximum of two disks restored in parallel per VM. Restore the VMs to VM names that are appended with the time and date of the restore operation.

```
dsmc restore vm vm* -vmmaxrestoreparallelvm=4 -vmmaxrestoresessions=20  
-vmmaxrestoreparalleldisks=2 -vmname="*_<timestamp>"  
-targetpath=c:\hypervrestore
```

Task

Of all the VMs that match the pattern `vm*` except for `vm5`, restore four VMs in parallel to the `c:\hypervrestore` folder. In the restore operation, allocate five sessions per VM that is being restored, with a maximum of two disks restored in parallel per VM. Restore the VMs to VM names that are appended with the time and date of the restore operation.

```
dsmc restore vm vm=vm*;-vm=vm5 -vmmaxrestoreparallelvm=4  
-vmmaxrestoresessions=20 -vmmaxrestoreparalleldisks=2  
-vmname="*_<timestamp>" -targetpath=c:\hypervrestore
```

Related reference

[“Vmmaxrestoreparalleldisks” on page 196](#)

Specify the `vmmaxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

[“Vmmaxrestoreparallelvms” on page 197](#)

The `vmmaxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

[“Vmmaxrestoresessions” on page 202](#)

The `vmmaxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

Preview virtual machine restore operations

You can use the `-preview` parameter to verify the results of a restore operation without restoring virtual machines (VMs). The `-preview` parameter provides a list of VMs that will be restored and information about the VMs. To understand how to use the `-preview` parameter with the **restore vm** command, review information about the options that are displayed and examples of the **restore vm -preview** command.

The `-preview` parameter returns options and their values only if the options override the default values or if no default exists.

When you issue the `-preview` parameter with the **restore vm** command, the restore operation does not start. The restore operation starts only if the `-preview` parameter is removed from the command.

You can use the `-preview` parameter with any valid **restore vm** command.

Example

Task

Preview the operation to restore all VMs to new VMs in the `d:\vmstorage` directory. The new VM names are composed of the original VM name, which is appended with the time stamp of the restore.

Issue the following command:

```
dsmc restore vm * -vmname="*_<timestamp>" -targetpath=d:\vmstorage\  
-preview
```

Command output:

```
IBM Storage Protect  
Command Line Backup-Archive Client Interface  
  Client Version 8, Release 1, Level 7.0  
  Client date/time: 10/02/2018 14:20:25  
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.  
  
Node Name: HYPERVTEST  
Session established with server ASANA2: Linux/x86_64  
  Server Version 7, Release 1, Level 9.000  
  Server date/time: 10/02/2018 14:18:30  Last access: 10/02/2018 14:05:26  
  
Restore function invoked.  
  
Restore VM command started.  Total number of virtual machines to process: 3  
1.  Virtual machine name: 'HEISENBERG.RESTORED'  
    Mode: 'Incremental Forever - Incremental'  
    Backup time: IFINCR 07/20/2018 06:59:16  
  
        Disk 1 label:          'Hard Disk 1'  
        Disk 1 name:          'IDE 0 0|HEISENBERG (Windows 2012 R2) Disk 1.vhdx'  
        Disk 1 status:        Selected  
        Disk 1 capacity:      136,365,211,648  
        Disk 1 data to send:  9,965,666,304  
  
        Destination name:     'HEISENBERG.RESTORED_10-02-2018_14-20-34'  
        Target Path:          'd:\vmstorage\  
  
2.  Virtual machine name: 'fremont1vm8'  
    Mode: 'Incremental Forever - Full'  
    Backup time: IFFULL 07/31/2018 11:33:43  
  
        Disk 1 label:          'Hard Disk 1'  
        Disk 1 name:          'IDE 0 0|fremont1vm8.vhdx'
```

```

Disk 1 status:      Selected
Disk 1 capacity:    42,949,672,960
Disk 1 data to send: 8,426,356,736

Destination name:    'fremont1vm8_10-02-2018_14-20-34'
Target Path:         'd:\vmstorage\'

```

```

3.    Virtual machine name: 'HEISENBERG (Windows 2012 R2)'
      Mode: 'Incremental Forever - Incremental'
      Backup time: IFINCR 09/17/2018 09:39:01

```

```

Disk 1 label:        'Hard Disk 1'
Disk 1 name:          'IDE 0 0|HEISENBERG (Windows 2012 R2) Disk 1.vhdx'
Disk 1 status:        Selected
Disk 1 capacity:      136,365,211,648
Disk 1 data to send:  9,730,785,280

```

```

Disk 2 label:        'Hard Disk 2'
Disk 2 name:          'IDE 1 1|Data Disk 1.vhdx'
Disk 2 status:        Selected
Disk 2 capacity:      2,147,483,648
Disk 2 data to send:  134,217,728

```

```

Disk 3 label:        'Hard Disk 3'
Disk 3 name:          'SCSI 0 0|Data Disk 2.vhdx'
Disk 3 status:        Selected
Disk 3 capacity:      1,073,741,824
Disk 3 data to send:  100,663,296

```

```

Destination name:    'HEISENBERG (Windows 2012 R2)_10-02-2018_14-20-34'
Target Path:         'd:\vmstorage\'

```

Task

Preview the operation to restore all VMs except for the VM named "fremont1vm8" to new VMs in the d:\vmstorage directory. The new VM names are composed of the original VM name, which is appended with the time stamp of the restore.

Issue the following command:

```

dsmc restore vm vm=*;-vm=fremont1vm8 -vmname="*_<timestamp>"
-targetpath=d:\vmstorage\ -preview

```

Command output:

```

IBM Storage Protect
Command Line Backup-Archive Client Interface
  Client Version 8, Release 1, Level 7.0
  Client date/time: 10/02/2018 14:24:13
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.

```

```

Node Name: HYPERVTEST
Session established with server ASANA2: Linux/x86_64
  Server Version 7, Release 1, Level 9.000
  Server date/time: 10/02/2018 14:22:17  Last access: 10/02/2018 14:18:37

```

Restore function invoked.

Restore VM command started. Total number of virtual machines to process: 2

```

1.    Virtual machine name: 'HEISENBERG.RESTORED'
      Mode: 'Incremental Forever - Incremental'
      Backup time: IFINCR 07/20/2018 06:59:16

```

```

Disk 1 label:        'Hard Disk 1'
Disk 1 name:          'IDE 0 0|HEISENBERG (Windows 2012 R2) Disk 1.vhdx'
Disk 1 status:        Selected
Disk 1 capacity:      136,365,211,648
Disk 1 data to send:  9,965,666,304

```

```

Destination name:      'HEISENBERG.RESTORED_10-02-2018_14-24-21'
Target Path:           'd:\vmstorage\'

2.    Virtual machine name: 'HEISENBERG (Windows 2012 R2)'
      Mode: 'Incremental Forever - Incremental'
      Backup time: IFINCR 09/17/2018 09:39:01

      Disk 1 label:       'Hard Disk 1'
      Disk 1 name:        'IDE 0 0|HEISENBERG (Windows 2012 R2) Disk 1.vhdx'
      Disk 1 status:      Selected
      Disk 1 capacity:    136,365,211,648
      Disk 1 data to send: 9,730,785,280

      Disk 2 label:       'Hard Disk 2'
      Disk 2 name:        'IDE 1 1|Data Disk 1.vhdx'
      Disk 2 status:      Selected
      Disk 2 capacity:    2,147,483,648
      Disk 2 data to send: 134,217,728

      Disk 3 label:       'Hard Disk 3'
      Disk 3 name:        'SCSI 0 0|Data Disk 2.vhdx'
      Disk 3 status:      Selected
      Disk 3 capacity:    1,073,741,824
      Disk 3 data to send: 100,663,296

      Destination name:   'HEISENBERG (Windows 2012 R2)_10-02-2018_14-24-21'
      Target Path:        'd:\vmstorage\'

```

Related reference

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

Chapter 10. Options reference

The following sections contain detailed information about each of the client options that are used for IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V operations.

Information for each option includes the following information:

- a description
- a syntax diagram
- detailed descriptions of the parameters
- examples of using the option in the client options file (if applicable)
- examples of using the option on the command line (if applicable)

Options with a command-line example of **Does not apply** cannot be used with command line or scheduled commands.

Dateformat

The `dateformat` option specifies the format you want to use to display or enter dates.

Use this option if you want to change the default date format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time you start the client. Consult the documentation on your local system for details about setting up your locale definition.

You can use the `dateformat` option with the **expire** command.

When you include the `dateformat` option with a command, it must precede the `fromdate` and `pitdate` options.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Date Format** drop-down list of the Preferences editor.

Syntax

➤ DATEformat — — *format_number* ➤

Parameters

format_number

Displays the date using one of the following formats. Select the number that corresponds to the date format you want to use:

1

MM/DD/YYYY

This is the default for the following available translations:

- US English
- Chinese (Traditional)
- Korean

2

DD-MM-YYYY

This is the default for the following available translations:

- Brazilian Portuguese
- Italian

3

YYYY-MM-DD

This is the default for the following available translations:

- Japanese
- Chinese (Simplified)
- Polish

4

DD.MM.YYYY

This is the default for the following available translations:

- German
- French
- Spanish
- Czech
- Russian

5

YYYY.MM.DD

This is the default for the following available translations:

- Hungarian

6

YYYY/MM/DD

7

DD/MM/YYYY

Examples

Options file:

```
dateformat 3
```

Command line:

```
-date=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Additional considerations for specifying time and date formats

The date or time format you specify with this option must be used when using options that take date and time as input. Examples are: `totime`, `fromtime`, `today`, `fromdate`, and `pittime`.

For example, if you specify the `timeformat` option as `TIMEFORMAT 4`, the value that you provide on the `fromtime` or `totime` option must be specified as a time such as `12:24:00pm`. Specifying `13:24:00` would not be valid because `TIMEFORMAT 4` requires an hour integer that is 12 or less. If you want to specify up to 24 hour values on an option, and if you want to use commas as separators, you must specify `TIMEFORMAT 2`.

Detail

Use the `detail` option to display management class, file space, and backup information.

Use the `detail` with the **query vm** command to display the following statistics:

- The average number of IBM Storage Protect objects that are needed to describe a single megablock, across all megablocks in a backup.
- The average number of IBM Storage Protect objects that are needed to describe a single megablock, for all megablocks in a file space.
- The number of backups that were created since the last full backup was created from the production disks.

The values returned on **query vm** can help you fine tune the heuristics (see the `Mbobjrefreshthresh` and `Mbpctrefreshthresh` options) to fine tune the values trigger for megablock refreshes.

Syntax

➤ DETail ➤

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query vm -detail
```

Domain.vmfull

The `domain.vmfull` option specifies the virtual machines (VMs) to include in your full virtual machine image backup operations.

Domain.vmfull for Microsoft Hyper-V virtual machines

For Hyper-V VM backups, use the `domain.vmfull` option to specify which Hyper-V VMs are processed when you run a **backup vm** command, without specifying any Hyper-V VM names.

You can specify which VMs to process by using any of the following techniques:

- Use the `VM=` option and specify the name of a virtual machine.
- Provide a comma-separated list of virtual machine names.
- Use wildcard syntax to process virtual machines that match the name pattern.
- Use the `vmname:vhdX=` option to specify the VM hard disk (VHDX) to include or exclude during the Hyper-V backup operation of a VM.
- Use the `all-vm` domain-level parameter. You can also include one or more virtual machines by using the `VM=` keyword, or exclude VMs by using the `-VM=` syntax.

The virtual machines that are specified on the `domain.vmfull` option are processed only when the **backup vm** command is entered without specifying a virtual machine or a list of virtual machines on the command line.



Attention: For Microsoft Hyper-V operations, the only valid domain-level parameter for the `domain.vmfull` option is **all-vm**. You can also include one or more virtual machines by using the `VM=` keyword, or exclude virtual machines by using the `-VM=` syntax.

Supported Clients


This option can be used with supported Windows clients. This option can also be defined on the server.

Options file

Set this option in the client options, by using the command line, or by using the **VM Backup** tab of the Preferences editor.

Restriction: The `vmname:vhdv=vhdv_location` parameter cannot be set in the Preferences Editor. Include this setting in the options file, or on the command line when you run a **backup vm** command:

Syntax for Microsoft Hyper-V virtual machines



Syntax rules: Multiple keywords must be separated by a semicolon. There cannot be any spaces after the semicolons. Multiple machine or domain names must be separated by commas, with no space characters. For examples, see `vm=vmname`.

Parameters

vmname

Defines the virtual machine name that you want to process. You can supply a list of virtual machine host names by separating the names with commas (`vm1, VM2, Vm5`). The names are case-sensitive and must match the capitalization that is shown on the Hyper-V host in the **Hyper-V Manager > Virtual Machines** view.

vm=vmname

The `vm=` keyword specifies that the next set of values is a list of virtual machine names. The `vm=` keyword is the default and is not required.

In this example, `vm=` is not specified and commas are used to separate the machine names.

```
domain.vmfull my_vm1,my_vm2
```

If you specify multiple keywords, such as `vm=` and `-vm=`, the values that the keywords refer to must be separated by semicolons, with no intervening space characters:

```
domain.vmfull vm=my_vm1;vm=my_vm2
domain.vmfull -vm=my_vm3;-vm=my_vm4
```

Wildcard characters can be used to select virtual machine names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character, for example:

- Exclude all files that have "test" in the host name: `-vm=*test*`
- Include all virtual machines with names such as: "test20", "test25", "test29", "test2A":

```
vm=test2?
```

You can exclude a virtual machine from a backup operation by specifying the exclude operator (-) before the `vm=` keyword. For example, `-vm` is used to exclude a particular machine, or machines, from a domain level backup, such as, `ALL -VM`. If "vm1" is the name of a virtual machine, you can back up all of the virtual machines in the domain, but prevent the virtual machine "vm1" from being backed up. Set the following option:

```
domain.vmfull all-vm;-vm=vm1
```

You cannot use the exclude operator (-) to exclude a domain, such as `ALL -VM`. The exclude operator works only at the virtual machine name level.

vmname:vhdxd=vhdxd_location

This option specifies the location of the virtual machine hard disk (VHDX) to include in Hyper-V virtual machine backup operations.

The `vmname` variable specifies the name of the virtual machine to back up. Wildcard characters can be used to select virtual machine names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character.

The `:vhdxd=disk_location` keyword specifies the location of the virtual machine disk to include in the backup operation. The disk location specified by the `disk_location` variable must begin with "SCSI" or "IDE" followed by the controller number and device location number. For example:

```
domain.vmfull "vm1:VHDX=IDE 1 0"  
domain.vmfull "vm*:VHDX=SCSI 0 1"  
domain.vmfull "vm?:VHDX=SCSI 0 1"
```

You can exclude a virtual machine disk from backup operations by specifying the exclude operator (-) before the `vhdxd=` keyword. For example, use `-vhdxd=` to exclude a VM disk from the backup operation of a virtual machine. For example:

```
domain.vmfull "vm1:-VHDX=IDE 1 0"
```

If you specify multiple virtual machine disks to include or exclude, the `vhdxd=` or `-vhdxd=` keyword and associated values must be separated by colons, with no intervening space characters. For example:

```
domain.vmfull "vm1:vhdxd=IDE 1 0:vhdxd=SCSI 0 1"
```

If you specify multiple virtual machine names and virtual machine disks, the VM name and associated values must be separated by semicolons, with no intervening space characters. For example:

```
domain.vmfull "vm1:VHDX=IDE 1 0:VHDX=SCSI 0 1;vm2:VHDX=IDE 1 0:VHDX=SCSI 0 1"  
domain.vmfull "vm1:-VHDX=IDE 1 0:-VHDX=SCSI 0 1;vm2:-VHDX=IDE 1 0:-VHDX=SCSI 0 1"
```

all-vm

This option specifies that a **backup vm** operation processes all Hyper-V virtual machines that are known to the Hyper-V host.

Examples for Microsoft Hyper-V virtual machines

Options file:

Include all virtual machines in full VM backup operations.

```
domain.vmfull all-vm
```

Include all virtual machines in full VM backup operations, except for the ones that have a name suffix of `_test`.

```
domain.vmfull all-vm;-vm=*_test
```

Include all virtual machines in full VM backup operations, but exclude virtual machines `testvm1` and `testvm2`.

```
domain.vmfull all-vm;-VM=testvm1,testvm2
```

Include IDE disks (with controller 1 and disk location 0) and SCSI disks (with controller 0 and disk location 1) in Hyper-V backup operations of virtual machines vm1 and vm2.

```
domain.vmfull "vm1:VHDX=IDE 1 0:VHDX=SCSI 0 1;vm2:VHDX=IDE 1 0:VHDX=SCSI 0 1"
```

Restriction: You cannot use the `all-vm` option together with the `vmname:-vhdx=` option in a single domain specification in the options file or on the command line. For example, `domain1 = all-vm:-VHDX=SCSI 0 0` is not valid.

Related reference

[“Exclude.vmdisk” on page 172](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

[“Include.vmdisk” on page 176](#)

The `INCLUDE.VMDISK` option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

[“Vmprocessvmwithphysdisks” on page 206](#)

Use the `vmprocessvmwithphysdisks` option to control whether Hyper-V RCT virtual machine (VM) backups are processed if the VM has one or more physical disks (pass-through disks) provisioned.

Exclude.vmdisk

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

The `EXCLUDE.VMDISK` option specifies the label of a virtual machine's disk to be excluded from a **backup vm** operation. If you exclude a disk on the **backup vm** command, the command-line parameters override any `EXCLUDE.VMDISK` statements in the options file.

EXCLUDE.VMDISK for Microsoft Hyper-V virtual machines

Use the `EXCLUDE.VMDISK` option to exclude a virtual machine disk from Hyper-V backup operations.

Supported clients

This option can be used with all supported Windows clients.

Options file

Set this option in the client options file. Command-line parameters override statements in the options file.

Syntax

➤ `EXCLUDE.VMDISK — vmname disk_location` ➤

Parameters

vmname

Specifies the name of the VM that contains a disk that you want to exclude from a **backup vm** operation. The name is the virtual machine display name. You can specify only one VM name on each `EXCLUDE.VMDISK` statement. Specify additional `EXCLUDE.VMDISK` statements for each VM disk to exclude.

The VM name can contain an asterisk (*) to match any character string, and a question mark (?) to match any one character. If the VM name contains space characters, surround the name with quotation marks (" ").

Tip: If the VM name contains special characters, such as bracket characters ([) or (]), the VM name might not be correctly matched. If a VM name includes special characters, use a question mark (?) to represent the special characters.

For example, to exclude a SCSI virtual machine disk from the backup of a VM named "Windows VM3 [2012R2] ", use this syntax in the options file:

```
EXCLUDE.VMDISK "Windows VM3 ?2012R2?" "SCSI 0 1"
```

disk_location

Specify the location of the virtual machine hard disk (VHDX) to exclude from a Hyper-V backup operation. The disk location label must begin with "SCSI" or "IDE" followed by the controller number and device location number. Wildcard characters are not allowed.

Tip: Use the **backup vm** command with the **-preview** option to determine the location of disks in a given VM. See the "**Backup VM**" topic for the syntax.

Examples

Options file

Exclude the Windows system disk from all virtual machines that begin with WinVM in the following statement in the options file:

```
exclude.vmdisk WinVM* "IDE 0 0"
```

Virtual machine vm1 contains a virtual machine disk with IDE controller number 1 and device location 0. To exclude this virtual machine disk from **backup vm** operations, specify the following statement in the options file:

```
EXCLUDE.VMDISK vm1 "IDE 1 0"
```

Virtual machine vm2 contains a virtual machine disk with SCSI controller number 0 and device location 1. Exclude this disk from backup operations by specifying the following statement in the options file:

```
EXCLUDE.VMDISK vm2 "SCSI 0 1"
```

Command line

The command line examples show the use of the exclusion operator (-) before the vhdX= keyword, to indicate that the disk is to be excluded.

Exclude an IDE disk (with controller number 1 and device location 0) from the backup operation of virtual machine vm1:

```
dsmc backup vm "vm1:-vhdX=IDE 1 0"
```

Exclude a SCSI disk (with controller number 0 and device location 1) from the backup operation of virtual machine vm2:

```
dsmc backup vm "vm2:-vhdX=SCSI 0 1"
```

Restriction: You cannot use the **all -vm** option together with the **vmname:-vhdX=** option on the command line or in the options file.

Tips for restoring Hyper-V VMs with excluded disks

During a VM restore operation, an informational message is displayed to indicate that a VM disk is not restored because it was excluded during the backup operation. The restore operation also verifies whether the original disk file still exists in the restore destination folder. If the original disk file still exists, it is reconnected to the restored VM. Otherwise, an empty disk file is created with the same attributes (such as file name, disk size, and block size) and the empty disk file is connected to the restored VM.

After a restore operation, the controller and disk order in the restored VM remains the same as the original VM. You do not have to adjust the disk location in the **EXCLUDE.VMDISK** option for future backup operations of the restored VM.

However, if you remove a SCSI controller manually, all subsequent SCSI controllers' numbers are changed. For example, if you remove "SCSI 0", the next SCSI controller (previously "SCSI 1") becomes "SCSI 0". In this case, you must update the VM disk location in the `EXCLUDE.VMDISK` option.

The disk location information such as "SCSI 0 0" is displayed in messages for backup, restore, and query operations.

Related reference

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Domain.vmfull” on page 169](#)

The `domain.vmfull` option specifies the virtual machines (VMs) to include in your full virtual machine image backup operations.

[“Include.vmdisk” on page 176](#)

The `INCLUDE.VMDISK` option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

Inactive

Use the `inactive` option to display both active and inactive objects.

You can use the `inactive` option with the **query vm** and **restore vm** commands.

Important: When using the `inactive` option during a restore operation, also use the `pick` option because all versions are restored in an indeterminate order. This option is implicit when `pitdate` is used.

Syntax

➤ INActive ➤

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore vm VM1 -inactive
```

Include.vm


This option overrides the management class that is specified on the `vmmc` option.

The management class specified on the `vmmc` option applies to all backups. You can use the `include.vm` option to override that management class, for one or more virtual machines. The `include.vm` option does not override or affect the management class that is specified by the `vmctlmc` option. The `vmctlmc` option binds backed-up virtual machine control files to a specific management class.

Options File

Set this option in the client options file.

Syntax



Parameters

vmname

Required parameter. Specifies the name of a virtual machine that you want to bind to the specified management class. Only one virtual machine can be specified on each `include .vm` statement. However, you can specify as many `include .vm` statements as needed to bind each virtual machine to a specific management class.

You can include wildcards in the virtual machine name. An asterisk (*) matches any character string. A question mark (?) matches a single character. If the virtual machine name contains a space character, enclose the name in double quotation marks (").

Tip: If the virtual machine name contains special characters, type the question mark wildcard in place of the special characters when you specify the virtual machine name.

mgmtclassname

Optional parameter. Specifies the management class to use when the specified virtual machine is backed up. If this parameter is not specified, the management class defaults to the global virtual machine management class that is specified by the `vmmc` option.

Examples

Assume that the following management classes exist and are active on the IBM Storage Protect server:

- MCFORTESTVMS
- MCFORPRODVMS
- MCUNIQUEVM

Example 1

The following `include .vm` statement in the client options file binds all virtual machines that have names that begin with `VMTEST` to the management class called `MCFORTESTVMS`:

```
include.vm vmtest* MCFORTESTVMS
```

Example 2

The following `include .vm` statement in the client options file binds a virtual machine that is named `WHOPPER VM1 [PRODUCTION]` to the management class called `MCFORPRODVMS`:

```
include.vm "WHOPPER VM1 ?PRODUCTION?" MCFORPRODVMS
```

The virtual machine name must be enclosed in quotation marks because it contains space characters. Also, the question mark wildcard is used to match the special characters in the virtual machine name.

Example 3

The following `include .vm` statement in the client options file binds a virtual machine that is named `VM1` to a management class that is named `MCUNIQUEVM`:

```
include.vm VM1 MCUNIQUEVM
```

Related reference

[“Vmmc” on page 206](#)

Use the `vmmc` option to store virtual machine backups by using a management class other than the default management class.

Include.vmdisk

The `INCLUDE.VMDISK` option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

The `INCLUDE.VMDISK` option specifies the label of a VM disk to be included in a **backup vm** operation. If you include a disk on the **backup vm** command, the command-line parameters override any `INCLUDE.VMDISK` statements in the options file.

INCLUDE.VMDISK for Microsoft Hyper-V virtual machines

Use the `INCLUDE.VMDISK` option to include a VM disk from Hyper-V backup operations.

Supported clients

This option can be used with all supported Windows clients.

Options file

Set this option in the client options file. Command-line parameters override statements in the options file.

Syntax

➤ `INCLUDE.VMDISK` — *vmname disk_location* ➤

Parameters

vmname

Specifies the name of the VM that contains a disk that you want to include from a **backup vm** operation. The name is the virtual machine display name. You can specify only one VM name on each `INCLUDE.VMDISK` statement. Specify additional `INCLUDE.VMDISK` statements for each VM disk to include.

The VM name can contain an asterisk (*) to match any character string, and a question mark (?) to match any one character. If the VM name contains space characters, surround the VM name with quotation marks (" ").

Tip: If the VM name contains special characters, such as bracket characters ([] or { }), the VM name might not be correctly matched. If a VM name includes special characters, use a question mark (?) to represent the special characters.

For example, to include a SCSI VM disk in the backup of a virtual machine named "Windows VM3 [2012R2] ", use this syntax in the options file:

```
INCLUDE.VMDISK "Windows VM3 ?2012R2?" "SCSI 0 1"
```

disk_location

Specify the location of the VM disk to include in a Hyper-V backup operation. The disk location label must begin with "SCSI" or "IDE" followed by the controller number and device location number. Wildcard characters are not allowed.

Tip: Use the **backup vm** command with the `-preview` option to determine the location of disks in a given virtual machine. See the "**Backup VM**" topic for the syntax.

Examples

Options file

Virtual machine vm1 contains an IDE VM disk (VHDX) at controller number 1 and device location 0. To include this VHDX in **backup vm** operations, specify the following statement in the options file:

```
INCLUDE.VMDISK vm1 "IDE 1 0"
```

Virtual machine vm2 contains a SCSI VM disk at controller number 0 and device location 1. Include this VHDX in backup operations by specifying the following statement in the options file:

```
INCLUDE.VMDISK vm2 "SCSI 0 1"
```

Command line

Include a single IDE disk (at controller number 1 and device location 0) when virtual machine vm1 is backed up:

```
dsmc backup vm "vm1:vhdX=IDE 1 0"
```

Include a SCSI disk (at controller number 0 and device location 1) in the backup operation of virtual machine vm2:

```
dsmc backup vm "vm2:vhdX=SCSI 0 1"
```

Related reference

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Domain.vmfull” on page 169](#)

The `domain.vmfull` option specifies the virtual machines (VMs) to include in your full virtual machine image backup operations.

[“Exclude.vmdisk” on page 172](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

INCLUDE.VMSNAPSHOTATTEMPTS

Use the `INCLUDE.VMSNAPSHOTATTEMPTS` option to determine the total number of snapshot attempts to try for a virtual machine (VM) backup operation that fails due to snapshot failure.

Supported Clients

This option can be used with supported Windows clients that are configured to back up VMs on Hyper-V hosts that run on Windows Server 2016 operating systems.

Options File

This option is valid in the client options file (`dsm.opt`). It can also be included on the server in a client option set. It is not valid on the command line.

Syntax

➤ `INCLUDE.VMSNAPSHOTATTEMPTS` — *vmname* — *num_with_quiescing* — *num_without_quiescing* ➤

Parameters

vmname

A required positional parameter that specifies the name of the virtual machine to attempt the total number of snapshots for, if a backup attempt fails due to snapshot failure. The name is the virtual machine display name.

Only one virtual machine can be specified on each `INCLUDE.VMSNAPSHOTATTEMPTS` statement. However, to configure the total snapshot attempts for other virtual machines, you can use the following methods:

- For each virtual machine that you want this option to apply to, specify as many `INCLUDE.VMSNAPSHOTATTEMPTS` statements as needed to reattempt snapshots that failed.
- Use wildcard characters for the *vmname* parameter value to specify virtual machine names that match the wildcard pattern. An asterisk (*) matches any character string. A question mark (?) matches a single character. If the virtual machine name contains a space character, enclose the name in double quotation marks ("").

Tip: If the virtual machine name contains special characters, type the question mark wildcard (?) in place of the special characters when you specify the virtual machine name.

num_with_quiescing

A positional parameter that specifies the following action:

For Hyper-V RCT backup operations:

The *num_with_quiescing* parameter specifies the number of times to attempt snapshots with quiescing to create application-consistent backups.

You can specify a value in the range 0 - 10. The default value is 2.

num_without_quiescing

For Hyper-V RCT backup operations:

The *num_without_quiescing* option specifies the number of times to attempt snapshots without quiescing after the specified number of attempts in the *num_with_quiescing* option are completed.

You can specify a value in the range 0 - 10. The default value is 0.

Important: When this parameter is applied to a VM backup, the backup is considered crash-consistent. As a result, operating system, file system, and application consistency are not guaranteed. An `include.vmsnapshotattempts 0 0` entry is not valid. Backup operations require at least one snapshot.

Examples

Hyper-V examples:

Example 1

Specify the following statement in the client options file to make two total snapshot attempts at crash-consistent backups for all Hyper-V VMs that begin with `LinuxVM`:

```
INCLUDE.VMSNAPSHOTATTEMPTS LinuxVM* 0 2
```

Example 2

Specify the following statement in the client options file to try three snapshot attempts for virtual machine `VM1`: two application-consistent snapshot attempts, and if they fail, to try one crash-consistent snapshot attempt:

```
INCLUDE.VMSNAPSHOTATTEMPTS VM1 2 1
```

INCLUDE.VMTSMVSS

Use the INCLUDE.VMTSMVSS option to enable application protection during backup operations of guest virtual machines (VMs) that host application data.

The INCLUDE.VMTSMVSS option notifies applications on the guest VM that a backup is about to occur. This option allows the application to truncate transaction logs and commit transactions so that the application can resume from a consistent state when the backup completes. An optional parameter can be specified to suppress the truncation of Microsoft SQL Server transaction logs.

When a VM is included by this option, application protection is provided. That is, the data mover freezes and thaws the VSS writers and, optionally, truncates the application logs. If a VM is not protected by this option, application protection is provided by Hyper-V, which freezes and thaws the VSS writers, but does not truncate application logs.

Important: Before you begin application protection backups, ensure that the application database, such as the Microsoft SQL Server database or Microsoft Exchange Server database, is on a non-boot drive (any drive other than the boot drive), in case a **diskshadow revert** operation is needed during restore.

Options file

Set this option in the data mover options file. This option cannot be set by the preferences editor or on the command line.

Syntax

➤ INCLUDE.VMTSMVSS — *vmname* — — OPTions=KEEPSqllog ➤

Parameters

vmname

Specifies the name of the VM that contains the applications to quiesce. The name is the VM display name in the Hyper-V Manager. Specify one VM per INCLUDE.VMTSMVSS statement. For example, to include a VM named Windows VM3 [2012R2], use the following syntax in the options file:

```
INCLUDE.VMTSMVSS "Windows VM3 [2012R2]"
```

To protect all VMs with this option, use an asterisk as a wildcard (INCLUDE.VMTSMVSS *). You can also use question marks to match any single character. For example, INCLUDE.VMTSMVSS vm?? protects all VMs that have names that begin with vm and are followed by any two characters (vm10, vm11, vm17, and so on).

Tip: If the VM name contains special characters, such as bracket characters ([or]), the VM name might not be correctly matched. If a VM name contains special characters, you can use the question mark character (?) to match the special characters in the VM name.

There is no default value for this parameter. To enable application protection, you must include VMs to be protected on one or more INCLUDE.VMTSMVSS statements. Make sure that you do not exclude a disk on a VM (by using the EXCLUDE.VMDISK option) if the disk contains application data that you want protected.

OPTions=KEEPSqllog

For Microsoft SQL Server only: If the OPTions=KEEPSqllog parameter is specified on an INCLUDE.VMTSMVSS statement, the parameter prevents SQL server logs from being truncated when a data mover that is installed on a data mover node backs up a VM that is running a SQL server.

Specifying this parameter allows the SQL server administrator to manually backup, and possibly truncate the SQL server logs, so that they can be preserved and be used to restore SQL transactions to a specific checkpoint, after the VM is restored.

When this option is specified, the SQL log is not truncated and the following message is displayed and logged on the server:

ANS4179I IBM Storage Protect application protection did not truncate the Microsoft SQL Server logs on VM 'VM'.

You can remove the `OPTIONS=KEEPSQLLOG` option to enable truncation of the SQL logs when a backup completes.

Note: The client does not back up the SQL log files. The SQL administrator must back up the log files so that they can be applied after the database is restored.

Examples

Options file

Configure application protection for a VM that is named `vm_example`:

```
INCLUDE.VMTSMVSS vm_example
```

For SQL Server: Configure application protection for `vm11`, `vm12`, and `vm15`:

```
INCLUDE.VMTSMVSS vm11
INCLUDE.VMTSMVSS vm12
INCLUDE.VMTSMVSS vm15 options=keepsqlllog
```

Command line

Not applicable; this option cannot be specified on the command line.

Related concepts

[“Shadow copy considerations for restoring an application protection backup from the data mover” on page 180](#)

For Windows virtual machines (VMs), if you attempt to restore an application protection backup from the data mover, be aware of shadow copy restrictions when you restore the application protection backup.

Related reference

[Exclude.vmdisk](#)

The `EXCLUDE.VMDISK` option excludes a virtual machine disk from backup operations.

[Include.vmdisk](#)

The `INCLUDE.VMDISK` option includes a virtual machine (VM) disk in backup operations. If you do not specify one or more disk labels, all disks in the VM are backed up.

[“INCLUDE.VMSNAPSHOTATTEMPTS” on page 177](#)

Use the `INCLUDE.VMSNAPSHOTATTEMPTS` option to determine the total number of snapshot attempts to try for a virtual machine (VM) backup operation that fails due to snapshot failure.

Shadow copy considerations for restoring an application protection backup from the data mover

For Windows virtual machines (VMs), if you attempt to restore an application protection backup from the data mover, be aware of shadow copy restrictions when you restore the application protection backup.

The shadow storage might run out of space

If you attempt to run a full VM restore of an application protection backup, the system provider snapshot is present on the restored VM. As the application writes to the disk, the shadow storage space grows until it runs out of disk space.

In general, if application protection was used during a backup, use only application protection restore to restore a database. When you restore the application, the volume is automatically reverted. However, if you must restore the full VM, you must either manually revert or delete the shadow copy.

After you restore the entire VM, verify that the restore was successful, and the data is not corrupted. If the data is not corrupted, delete the shadow copy. If the data is corrupted, revert the shadow copy to restore data integrity.

You can determine which shadow copy to delete or revert by looking for the `dsmShadowCopyID.txt` file in the root directory of each restored volume. This file contains the snapshot IDs of the shadow copies that were created during the snapshot attempts. You can use the **diskshadow** command **delete shadows** to delete these IDs, or the **revert** command to revert the shadow copy. After the delete or revert is completed, you can also delete the `dsmShadowCopyID.txt` file.

Important: In order for the revert operation to succeed, the application database, such as the Microsoft SQL Server database or Microsoft Exchange Server database, must be on a non-boot drive (any drive other than the boot drive).

The shadow copy must be available on the restored volume during an application protection restore

In some cases, an application protection backup operation might use the Volume Shadow Copy Service (VSS) to create an application-consistent shadow copy before you start a VM backup. All changes that are made after the creation time of the shadow copy are saved to the shadow storage.

A database restore might fail if the shadow copy is not available during an application restore. The shadow copy is used at the time of restore to revert the restored volume to an application-consistent state. If the shadow copy is not available, the restored data will be in an inconsistent state.

The following situations can cause the shadow copy to be unavailable:

- Typically, the shadow storage is part of a volume. However, sometimes the shadow storage space is configured to be on a different volume either by default or manually. In this case, the database restore might fail because the shadow copy that was created during the VM backup operation is not available at restore time.
- The shadow storage is not available because the volume with the shadow storage was excluded at backup time.

The following workarounds are available for this issue:

- Before you run a VM backup, add the shadow copy storage association for each volume that is available on the guest VM by using the **vssadmin add shadowstorage** command. For example, to set the shadow storage location for volume E: on volume E:, issue the following command:

```
vssadmin add shadowstorage /for=E: /on=E: /maxsize=unbounded
```

Important: The **vssadmin add shadowstorage** command might fail if the VM has existing VSS snapshots. You must delete the VSS snapshots by using the same application that created them.

For example, if a VSS backup of an Exchange database with LOCAL backup destination was created by IBM Storage Protect for Mail: Data Protection for Microsoft Exchange Server, use the Data Protection for Microsoft Exchange Server application to delete the VSS backup. If an unidentified VSS snapshot exists, use the Windows **diskshadow** command **delete shadows** to delete the VSS snapshot.

Also, ensure that the volume that holds the shadow storage is not excluded from backup operations.

- Manually revert snapshots to achieve application-consistency of the database files:
 1. Mount all disks in the VM backup by using IBM Storage Protect recovery agent.
 2. Start the Windows **diskshadow** command in interactive mode.
 3. In the interactive **diskshadow** mode, issue the following command:

```
list shadows all
```

4. In the root directory of each mounted drive, locate the `dsmShadowCopyID.txt` file. This file contains the globally unique identifier (GUID) of the VSS shadow copy that is needed in the volume revert operation.
5. Open the `dsmShadowCopyID.txt` file and identify the GUID of the volume where the database files are located.

6. In the interactive **diskshadow** mode, issue the following command:

```
revert GUID
```

where *GUID* is the snapshot GUID that was identified in the `dsmShadowCopyID.txt` file.

In order for the revert operation to succeed, the application database must be on a non-boot drive.


Mode

Use the mode option to specify the backup mode to use when performing specific backup operations.

You can use the mode option with the **backup vm** command. this parameter specifies whether to perform a full image backup, an incremental-forever full backup, or an incremental-forever-incremental backup of Hyper-V virtual machines.

The mode option has no effect on a when backing up a raw logical device.

Syntax



Parameters

IFIncremental

Specifies that you want to perform an incremental-forever-incremental backup of a Hyper-V virtual machine. An incremental-forever-incremental backup backs up only the disk blocks that have changed since the last backup. This is the default backup mode.

You cannot use this backup mode to back up a virtual machine if the client is configured to encrypt the backup data.

IFFull

Specifies that you want to perform an incremental-forever-full backup of a Hyper-V virtual machine. An incremental-forever-full backup backs up all used blocks on a virtual machine's disks. By default, the first backup of a Hyper-V virtual machine is an incremental-forever-full (mode=iffull) backup, even if you specify mode=ifincremental (or let the mode option default). Subsequent backups default to mode=ifincremental.

You cannot use this backup mode to back up a virtual machine if the client is configured to encrypt the backup data.

Examples

Task

Perform an incremental-forever-full VM backup of a Windows Hyper-V VM named `msvm1`

```
dsmc backup vm msvm1 -mode=iffull
```

Task

Perform an incremental-forever-incremental backup of a Windows Hyper-V VM named `msvm1`

```
dsmc backup vm msvm1 -mode=ifincremental
```

Related reference

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

Mbobjrefreshthresh

The `mbobjrefreshthresh` (megablock object refresh threshold) option is a number defining a threshold. When the number of IBM Storage Protect objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

When you backup a virtual machine, the data is stored on the IBM Storage Protect server in 128 MB units, called *megablocks*. If an area on the production disk changes and a new incremental backup is performed, a new megablock is created to represent the changes that were made to the previously backed up data. Because a new megablock can be created with each incremental backup, eventually the megablocks can adversely affect the performance of the IBM Storage Protect database, and therefore, adversely affect the performance of most IBM Storage Protect operations.

Use this option when estimating IBM Storage Protect objects that represent production data for each virtual machine backup. For example, when the number of IBM Storage Protect objects exceed this value, the megablock is refreshed. This action means that the entire 128-MB block is copied to the IBM Storage Protect server and is represented as a single IBM Storage Protect object. The minimum value is 2 and the maximum value is 8192. The default value is 50.

Options file

This option is valid in the client options file (`dsm.opt`). It can also be included on the server in a client option set. It is not valid on the command line.

Syntax



Parameters

The minimum value you can specify is 2 megablocks, the largest value is 8192 megablocks; the default is 50 megablocks.

Examples

Set this option to trigger a megablock refresh when the number of objects needed to represent an updated megablock exceeds 20 objects:

```
MBOBJREFRESHTHRESH 20
```

Mbpctrefreshthresh

The `mbpctrefreshthresh` (megablock percentage refresh threshold) option is a number defining a threshold. When the number of IBM Storage Protect percentage of objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

When you backup a virtual machine, data is stored on the IBM Storage Protect server in 128 MB units, called *megablocks*. If an area on the production disk changes and a new incremental backup is performed, a new megablock is created to represent the changes that were made to the previously backed up data. Because a new megablock can be created with each incremental backup, eventually the megablocks can adversely affect the performance of the IBM Storage Protect database, and therefore, adversely affect the performance of most IBM Storage Protect operations.

Use this option when estimating the amount of additional data that is backed up for each virtual machine. For example, when a 128-MB block of a production disk changes more than the percentage specified, the entire 128-MB block is copied to the IBM Storage Protect server. The block is represented as a single IBM Storage Protect object.

Options file

This option is valid in the client options file (`dsm.opt`). It can also be included on the server in a client option set. It is not valid on the command line.

Syntax



Parameters

The minimum value you can specify is 1 percent, the largest value is 99 percent; the default is 50 percent.

Examples

Set this option to trigger a megablock refresh when 50 percent (or more) of the objects in a megablock on a production disk have changed:

```
MBPCTREFRESHTHRESHOLD 50
```

Noprompt

The `noprompt` option suppresses the confirmation prompt that is presented by the **expire** command. Use the `noprompt` option with the **expire** command.

Syntax

➤ `NOPrompt` ➤

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc expire -noprompt c:\home\project\*
```

Numberformat

The `numberformat` option specifies the format you want to use to display numbers.

Use this option if you want to change the default number format for the language of the message repository you are using.

By default, format information is obtained from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

You can only use the `numberformat` option with the **expire** command.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Number Format** field of the Preferences editor.

Syntax

➤ NUMberformat — — *number* ➤

Parameters

number

Displays numbers using any one of the following formats. Specify the number (0–6) that corresponds to the number format you want to use.

0

Use the locale-specified date format. This is the default (does not apply to Mac OS X).

1

1,000.00

This is the default for the following available translations:

- US English
- Japanese
- Chinese (Traditional)
- Chinese (Simplified)
- Korean

2

1,000,00

3

1 000,00

This is the default for the following available translations:

- French
- Czech
- Hungarian
- Polish
- Russian

4

1 000.00

5

1.000,00

This is the default for the following available translations:

- Brazilian Portuguese
- German
- Italian
- Spanish

6

1'000,00

Examples

Options file:

```
num 4
```

Command line:

```
-numberformat=4
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Pick

The pick option creates a list of backup versions or archive copies that match the file specification you enter.

From the list, you can select the versions to process. Include the `inactive` option to view both active and inactive objects.

Use the pick option with the **restore vm** command.

Syntax

➤ Pick ➤

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore vm vmfin* -pick -inactive
```

Pitdate

Use the pitdate option with the pittime option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored.

Use the pitdate option with the **query vm** and **restore vm** commands.

When pitdate is used, the `inactive` and `latest` options are implicit.

Syntax

➤ PITDate = — — date ➤

Parameters

date

Specifies the appropriate date.

Examples

Command line:

```
dsmc restore vm vmfin3 -pitdate=02/21/2014
```

Pitttime

Use the **pitttime** option with the **pitdate** option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify **pitdate** option.

Use the **pitttime** option with the **query vm** and **restore vm** commands.

Syntax

➡ PITTime = — — *time* ➡

Parameters

time

Specifies a time on a specified date. If you do not specify a time, the time defaults to 23:59:59.

Examples

Command line:

```
dsmc query vm vmfin1 -pitt=06:00:00 -pitd=02/03/2014
```

Skipsystemexclude

Use the **skipsystemexclude** option to specify how to process exclude statements for certain operating system files that the IBM Storage Protect for Virtual Environments client skips by default.

By default, IBM Storage Protect for Virtual Environments clients skip certain Windows operating system files that are not normally required for system recovery during virtual machine (VM) backup operations. These files can include Windows system files, temporary internet files, and files in the Recycle Bin.

You can use this option to skip the processing of exclude statements for these operating system files. By not processing these exclude statements, the time it takes to back up VMs might be reduced.

Support clients

This option is valid for IBM Storage Protect for Virtual Environments clients on Windows operating systems only.

Options file

This option is valid in the client options file (**dsm.opt**) or on the command line.

Syntax

➡ SKIPSYSTemexclude — { Yes No } ➡

Parameters

Yes

Specify this parameter to skip the processing of exclude statements for certain Windows operating system files during VM backup operations. This parameter is the default.

No

Specify this parameter to process exclude statements of Windows operating system files. When you select this parameter and run a file backup of the Hyper-V host, the operating system files are excluded.

Examples

Options file

```
SKIPSYSTemexclude yes
```

Command line

```
dsmc backup vm -SKIPSYST=yes
```

```
dsmc incr -skipsyst=no
```

Timeformat

The `timeformat` option specifies the format in which you want to display and enter system time.

Use this option if you want to change the default time format for the language of the message repository you are using.

By default, format information is obtained from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

You can only use the `timeformat` option with the **expire** command.

When you include the `timeformat` option with a command, it must precede the `fromtime`, `pittime`, and `totime` options.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Time Format** field of the Preferences editor.

Syntax

►► TIMEformat — — *format_number* ◄◄

Parameters

format_number

Displays time in one of the formats listed here. Select the format number that corresponds to the format you want to use. When you include the `timeformat` option in a command, it must precede the `pittime` option.

1

23:00:00

2

23,00,00

3

23.00.00

4

12:00:00 A/P

5

A/P 12:00:00

Examples

Options file:

```
timeformat 4
```

Command line:

```
-time=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Additional considerations for specifying time and date formats

The date or time format you specify with this option must be used when using options that take date and time as input. Examples are: `totime`, `fromtime`, `today`, `fromdate`, and `pittime`.

For example, if you specify the `timeformat` option as `TIMEFORMAT 4`, the value that you provide on the `fromtime` or `totime` option must be specified as a time such as `12:24:00pm`. Specifying `13:24:00` would not be valid because `TIMEFORMAT 4` requires an hour integer that is 12 or less. If you want to specify up to 24 hour values on an option, and if you want to use commas as separators, you must specify `TIMEFORMAT 2`.

Vmautostartvm

Use the `vmautostartvm` option to specify whether the virtual machine (VM) that is created during instant access processing is automatically powered on. This option is valid only when used with the **restore VM** `vmrestoretype=instantaccess` command.


Supported Clients

This option can be used with supported Windows clients.

Options file

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

NO

The VM that is created for instant access is not started automatically. The VM must be started manually. This parameter is the default setting. The default provides an opportunity to reconfigure the VM before you power it on to avoid potential conflicts with existing VMs.

YES

The VM that is created for instant access is started automatically.

Examples

Options file:

```
VMAUTOSTARTvm NO
```

Command line:

```
dsmc restore vm Oslo -VMRESToretype=INSTANTAccess -vmname=Oslo_verify  
-VMAUTOSTARTvm=YES
```

Vmbackdir

The **vmbackdir** option specifies the temporary disk location where the client saves control files that are created during full VM backup and restore operations of Microsoft Hyper-V virtual machines.

When a client on a data mover node starts a full VM backup of a virtual machine, the client creates metadata in files that are associated with the backed up virtual machine and its data. The files that contain the metadata are referred to as *control files*.

During full VM backup operations, the metadata is saved on a disk in the data mover node until the backup completes and both the virtual machine data and the control files are saved to server storage. During a full VM restore operation, the control files are copied from the server and are temporarily stored on the data mover disk, where they are used to restore the virtual machine and its data. After a backup or a restore operation completes, the control files are no longer needed and the client deletes them from their temporary disk location.

The directory that is specified by this option must be on a drive that contains sufficient free space to contain the control information from a full VM backup.

Options File

Set this option in the client options file, or specify it on the command line as an option for the **backup vm** or **restore vm** commands.

Syntax

➤ VMBACKDir — directory ➤

Parameters

directory

Specifies the path where the control files are stored on the backup server.

The default is c:\mnt\tsmvmbackup\fullvm\

Examples

Options file:

```
VMBACKD c:\mnt\tsmvmbackup\
```

Command line:

```
dsmc backup vm -VMBACKUPT=fullvm -VMBACKD=G:\virtual_machine\control_files\  
dsmc restore vm -VMBACKUPT=fullvm -VMBACKD=G:\san_temp\
```

Vmctlmc

This option specifies the management class to use when backing up virtual machine control files.

By default, virtual machine control files are bound to the default management class. The **vmmc** option can be used to specify a different management class to which virtual machine data and virtual machine control files are bound. The **vmctlmc** option overrides the default management class and the **vmmc** option for the virtual machine control files.

Under certain conditions, it might be desirable or necessary to bind the control files to a different management class than the data files.

The `vmctlmc` option is required if virtual machine data files are backed up to tape. Virtual machine control files must be backed up to a disk-based storage pool that does not migrate to tape. The storage pool can be composed of random access volumes and sequential file volumes; the storage pool can also be a deduplicated pool. Use the `vmctlmc` option to specify a management class that stores data in such a storage pool.

Restriction: The management class that is specified by the `vmctlmc` option determines only the destination storage pool for virtual machine control files. Retention of the control files is determined by the `vmmc` option, if specified, or by the default management class. The retention for the virtual machine control files always matches the retention of the virtual machine data files.

Options File

Place this option in the client options file `dsm.opt`.

Syntax

➤ `VMCTLmc` — *class_name* ➤

Parameters

class_name

Specifies a management class that applies to backing up virtual machine control files. If you do not set this option, the management class that is specified on the `vmmc` option is used. If you do not set this option and the `vmmc` option is not set, the default management class of the node is used.

Examples

Options file:

`vmctlmc diskonlymc`

Command line:

Does not apply.

Vmiscsiserveraddress

If there are multiple NICs for a Hyper-V host, use the `vmiscsiserveraddress` option to specify the IP address of the Hyper-V host server that is used for Linux File Restore in the same subnet as the iSCSI initiator and iSCSI target.

The `vmiscsiserveraddress` option is valid for all instant operations (`vmrestoretype=instantaccess` and `vmrestoretype=instantrestore`) for Hyper-V virtual machines.

Supported clients

This option is valid for supported Windows and Linux clients that operate as data movers to back up Hyper-V virtual machines.

Options file

If Hyper-V host has multiple NICs, for Linux File restore, you must set the `VMISCSIServeraddress` host IP address that will be used for Linux File Restore in the `dsm.hostname_HV_MP_WIN.opt` file.

Syntax

➤ `VMISCSIServeraddress` — — *Hyper-V_host_IP_Address* ➤

Parameters

Hyper-V_host_IP_Address

Specify the Hyper-V host IP address that will be used for Linux File restore. This IP address must be in the same subnet as the iSCSI initiator and iSCSI target. If `vmiscsiserveraddress` is not specified and if the Hyper-V has multiple NICs, a random IP address for the Hyper-V host will be selected for Linux File restore.

Examples

Options file:

```
VMISCSIServeraddress 192.168.42.50
```

Vmmaxbackupsessions

The `vmmaxbackupsessions` option specifies the maximum number of IBM Storage Protect server sessions that send virtual machine (VM) data to the server that can be included in an optimized backup operation.

An optimized backup operation is one in which parallel backup capability is enabled at the VM, virtual disk, or subdisk level.

The `vmmaxbackupsessions` option works with the `vmmaxparallel` option to optimize backup operations and to help control the amount of resources that the backup operation can create on a host in the Hyper-V environment. Adjust the values of these options to find the combination that provides optimal performance for the backup operations in your environment.


Supported clients

This option can be used with data movers on Windows Server 2012 R2, Windows Server 2016 and later operating systems. This option is not supported for data movers on Windows Server 2012.

Options file

This option is valid in the data mover options file (`dsm.hostname_HV_DM.opt`) or on the command line for use with the **backup vm** command. It can also be included in a client option set on the IBM Storage Protect server.

Syntax



Parameters

integer

Specifies the maximum number of IBM Storage Protect server sessions that can be created during the backup operation. The default is 4. The maximum is 100.

Tip: When you use the configuration wizard to configure Data Protection for Microsoft Hyper-V, the default is 16. This value is automatically updated in the data mover options file.

Review the following information when you use the `vmmaxbackupsessions` option along with the `vmmaxparallel` option or the `maxnummp` server parameter:

vmmaxparallel

The `vmmaxparallel` option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at any one time. 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, the following message is returned and the value is changed to the same value as the `vmmaxparallel` option:

```
ANS9995W The value of the VMMAxBACKUPSESSIONS option is number_value.  
This value must be greater than or equal to the value of the  
VMMAXPARALLEL option, which is number_value. The value will be  
set to the value of the VMMAXPARALLEL option.
```

maxnummp

The `maxnummp` server parameter specifies the maximum number of mount points that a node is allowed to use on the IBM Storage Protect server when the copy destination of the storage pool is `FILE` or `TAPE`. The `maxnummp` parameter setting must be equal to or greater than the `vmmaxparallel` and `vmmaxbackupsessions` option settings. When multiple instances of the data mover are backing up files, or when a single data mover runs parallel backups, more mount points might be needed.

If the value for the `vmmaxparallel` or `vmmaxbackupsessions` option exceeds the value for the `maxnummp` parameter, `ANS0266I` and other messages are displayed. Depending on the message, the data mover reduces the value of the `vmmaxparallel` option to match the number that is specified by the `maxnummp` parameter or prohibits additional sessions from being opened for the specified VM. In either situation, the backup operation continues.

If additional `ANS0266I` errors are detected, the data mover reduces the `vmmaxparallel` value by 1 and attempts to continue the backup operation. If `vmmaxparallel` is decremented to 1 and more `ANS0266I` errors are generated, the backup operation is ended and the following error is issued:

```
ANS5228E A backup VM operation failed because VMMAXPARALLEL was reduced  
to 1 and the client still cannot obtain a server mount point.
```

If you want to increase the value for the `maxnummp` parameter so your node can support additional parallel backup sessions, contact your server administrator.

Examples

Options file

```
VMMAxBACKUPS 10
```

Command line:

```
dsmc backup vm -VMMAxBACKUPS=10
```

Related tasks

[“Backing up multiple virtual machines in parallel \(optimized backup\)” on page 227](#)

With parallel backup processing, you can back up multiple virtual machines (VMs) at the same time to help optimize the performance of VM backup operations.

Related reference

[“Vmmaxparallel” on page 194](#)

The `vmmaxparallel` option is used to configure parallel backups of several virtual machines (VMs) by using a single instance of the data mover. This option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at a time.

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

[“Optimizing backup operations” on page 225](#)

To help optimize virtual machine (VM) backup operations, adjust the settings for the `vmmaxparallel` option. With this option, you can optimize backup operations without adversely affecting Hyper-V hosts.

Vmmaxparallel

The `vmmaxparallel` option is used to configure parallel backups of several virtual machines (VMs) by using a single instance of the data mover. This option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at a time.


An optimized backup operation is one in which parallel backup capability is enabled at the VM, virtual disk, or subdisk level.

The `vmmaxparallel` option works with the `vmmaxbackupsessions` option to optimize backup operations and to help control the amount of resources that the backup can create on a Hyper-V host. Adjust the values of these options to find the values that provide optimal performance for the backups that are in your environment.

Options file

This option is valid in the data mover options file (`dsm.hostname_HV_DM.opt`) or on the command line for use with the **backup vm** command. It can also be included on the server in a client option set.

Syntax



Parameters

integer

Specifies the maximum number of VMs that can be backed up at a time during a parallel backup operation. The default is 4. The maximum is 50.

Tip: When using client-side data deduplication, a deduplication session is started for each VM. This deduplication session is not counted as one of the `vmmaxparallel` sessions.

Review the following information when you use the `vmmaxparallel` option in conjunction with the `vmmaxbackupsessions` option or the `maxnummp` server parameter:

vmmaxbackupsessions

The `vmmaxbackupsessions` specifies the maximum number of sessions that send VM data to the IBM Storage Protect server that can be included in an optimized backup operation. The value of the `vmmaxbackupsessions` option must be equal to or greater than the value of the `vmmaxparallel` option.

maxnummp

The `maxnummp` server parameter specifies the maximum number of mount points that a node is allowed to use on the IBM Storage Protect server when the copy destination of the storage pool is FILE or TAPE. The `maxnummp` parameter setting must be equal to or greater than the `vmmaxparallel` and `vmmaxbackupsessions` option settings. When multiple instances of the data mover are backing up files, or when a single data mover runs parallel backups, more mount points might be needed.

If the value for the `vmmaxparallel` or `vmmaxbackupsessions` option exceeds the value for the `maxnummp` parameter, ANS0266I and other messages are displayed. Depending on the message, the data mover reduces the value of the `vmmaxparallel` option to match the number that is specified by the `maxnummp` parameter or prohibits additional sessions from being opened for the specified VM. In either situation, the backup operation continues.

If additional ANS0266I errors are detected, the data mover reduces the `vmmaxparallel` value by 1 and attempts to continue the backup operation. If `vmmaxparallel` is decremented to 1 and more ANS0266I errors are generated, the backup operation is ended and the following error is issued:

```
ANS5228E A backup VM operation failed because VMMAXPARALLEL was reduced
to 1 and the client still cannot obtain a server mount point.
```

If you want to increase the value for the `maxnummp` parameter so your node can support additional parallel backup sessions, contact your server administrator.

Examples

Options file

```
VMMAXP 10
```

Command line

```
dsmc backup vm -vmmaxp=10
```

Related reference

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

[“Domain.vmfull” on page 169](#)

The `domain.vmfull` option specifies the virtual machines (VMs) to include in your full virtual machine image backup operations.

[“Optimizing backup operations” on page 225](#)

To help optimize virtual machine (VM) backup operations, adjust the settings for the `vmmaxparallel` option. With this option, you can optimize backup operations without adversely affecting Hyper-V hosts.

Vmmaxpersnapshot

Use the `vmmaxpersnapshot` option to specify the maximum number of virtual machines (VMs) to include in a Hyper-V snapshot. The VMs in the snapshot are backed up to the IBM Storage Protect server.

By increasing the number of VMs in a snapshot, you can reduce the number of snapshots that are taken for a backup operation. This capability reduces the scheduling contention that can be experienced during cluster backup operations of VMs on Clustered Shared Volumes (CSVs).

A snapshot with more VMs takes longer to complete and increases the load on the system. A larger number of VMs means that the snapshot persists longer, which can affect performance.

This option is valid only for Hyper-V backup operations on Windows Server 2012 and 2012 R2 operating systems.

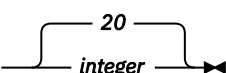
Supported clients

This option is valid for all supported Windows clients. This option can also be defined on the server.

Options file

This option is valid in the client options file (`dsm.opt`) or on the command line for the **Backup VM** command. It can also be included on the server in a client option set. It cannot be set in the Preferences Editor.

Syntax

➔ VMMAXPERSnapshot  ➔

Parameters

integer

Specifies the maximum number of VMs that can be included in a Hyper-V snapshot. The default is 20. The maximum is 100. The minimum is 1.

If some VMs reside on local volumes and some VMs reside on Clustered Shared Volumes (CSVs), the number of VMs in a snapshot might be less than the `vmmaxpersnapshot` setting. A snapshot cannot contain a mixture of VMs on local and CSV volumes.

To avoid creating a snapshot that spans volumes, the number of VMs in a snapshot might be less than the maximum number if the VMs are on different volumes. For example, four VMs are on Volume A and one VM is on Volume B. A snapshot is taken with only four VMs (from Volume A) even though the maximum setting is five. A second snapshot is taken for Volume B.

Examples

Options file

```
vmmaxpersnapshot 10
```

Command line

```
dsmc backup vm -vmmaxpers=10
```

Related concepts

[“Tuning scheduled VM backups for Windows Server 2012 and 2012 R2 clusters” on page 59](#)

Beginning with Data Protection for Microsoft Hyper-V version 8.1.2, you can back up more virtual machines (VMs) in parallel and across nodes in a cluster. A cluster node backup operation always retries the snapshot on volumes with snapshots that failed with a recoverable condition. You can also tune the number of VMs in a snapshot to reduce the workload of a snapshot for the Hyper-V host.

Related reference

[“Vmmaxsnapshotretry” on page 203](#)

Use the `vmmaxsnapshotretry` option to specify the maximum number of times to retry a snapshot operation of a virtual machine (VM) if the initial snapshot fails with a recoverable condition.

Vmmaxrestoreparalleldisks

Specify the `vmmaxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

You can specify the number of disk sessions to be opened, up to a maximum of 10. The maximum number of disk restore sessions is specified by the `vmmaxrestoresessions` option. The available sessions are allocated across the number of disk sessions as specified by the `vmmaxrestoreparalleldisks` option, by rounding down the number of sessions per disk to the nearest whole number.

Supported clients

This option can be used with supported Windows clients.

Options file

This option is valid in the data mover options file (`dsm.hostname_HV_DM.opt`) or on the command line with the **restore vm** command. It can also be included in a client option set on the IBM Storage Protect server.

Syntax



Parameters

integer

Specifies the number of virtual hard disks that can be restored simultaneously. The default is 2. The maximum is 10.

Example

Task

Set a maximum of 10 simultaneous restore operations for virtual disks in the restore operation of the VM **vm1**:

```
dsmc restore vm vm1 -vmmxrestoreparalleldisks=10 -vmmxrestoresessions=20
```

This command assigns two simultaneous restore sessions per virtual disk.

Related reference

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Vmmxrestoreparallelvms” on page 197](#)

The `vmmxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

[“Vmmxrestoresessions” on page 202](#)

The `vmmxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

[“Optimizing restore operations” on page 227](#)

To help optimize virtual machine (VM) restore operations, adjust the settings for the `vmmxrestoresessions`, `vmmxrestoreparallelvms`, and `vmmxrestoreparalleldisks` options. With these options, you can optimize restore operations while balancing the resources on the Hyper-V hosts.

Vmmxrestoreparallelvms

The `vmmxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

Use this option to help improve the performance of restore operations by increasing the number of VMs to restore in parallel.

You can specify the number of VMs to be restored simultaneously, up to a maximum of 10. The default value is 2.


Supported clients

This option can be used with supported Windows clients.

Options file

This option is valid in the data mover options file (`dsm.opt`) or on the command line for the **restore vm** command. It can also be included on the server in a client option set.

Syntax



Parameters

integer

Specifies the maximum number of VMs that can be restored simultaneously. The default is 2. The maximum is 10.

If you are using the `vmmaxrestoresessions` option to limit the number of restore sessions, the number of sessions must be greater than or equal to the number of VMs. This setting ensures that at least one session is available per VM.

If you are using the `vmmaxparalleldisks` option to restore multiple virtual disks at the same time, the number of virtual disks must be less than or equal to the number of sessions.

Examples

Task

Set a maximum of five simultaneous restore operations for VMs VM1, VM2, VM3, VM4, and VM5:

```
dsmc restore vm vm=VM1,VM2,VM3,VM4,VM5 -VMMAXRESTOREPARALLELVms=5
-VMMAXRESTORESessions=10 -VMMAXRESTOREPARALLELDisks=2
-asnode=kingston23_HV_TGT -optfile=dsm.KINGSTON23_HV_DM.opt -replace=yes
```

This command assigns up to 5 VM restore operations in parallel, and each VM can restore up to one virtual disk in parallel per VM at a time and assign 2 restore sessions per disk.

Command output:

```
IBM Storage Protect
Command Line Backup-Archive Client Interface
  Client Version 8, Release 1, Level 7.4
  Client date/time: 10/05/2018 15:01:01
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights
Reserved.

Node Name: KINGSTON23_HV_DM
Session established with server ASANA2: Linux/x86_64
  Server Version 7, Release 1, Level 9.000
  Server date/time: 10/05/2018 14:59:09  Last access: 10/05/2018 14:58:18

Accessing as node: KINGSTON23_HV_TGT
Restore function invoked.

Restore VM command started.  Total number of virtual machines to process: 5

Restore of Virtual Machine 'VM1' started

Restore of Virtual Machine 'VM2' started

Restore of Virtual Machine 'VM3' started

Restore of Virtual Machine 'VM4' started

Restore of Virtual Machine 'VM5' started

Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM1' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM1'
Restoring virtual disks for virtual machine 'VM1'
  Total disks to restore      : 1
  Disks to restore concurrently : 1
  Maximum Sessions per disk   : 2
```



```

Using 2 session(s) to restore virtual machine 'VM1' Disk 1 .
Restoring 'VM1' Hard Disk 1
  Type           : IFINCR
  Timestamp      : 10/05/2018 14:59:23
  Bytes to restore : 1,073,741,824

< 670.25 MB> [ - ]
Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM2' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM2'
Restoring 1,073,741,824 'VM1' e:\Hyper-V\VM1\Virtual Machines\1G.vhdx
[Done]
Restoring virtual disks for virtual machine 'VM2'
  Total disks to restore       : 1
  Disks to restore concurrently : 1
  Maximum Sessions per disk    : 2

Using 2 session(s) to restore virtual machine 'VM2' Disk 1 .
Restoring 'VM2' Hard Disk 1
  Type           : IFFULL
  Timestamp      : 10/05/2018 14:59:24
  Bytes to restore : 1,073,741,824

Restoring 0 VM1 [Done]
< 1.69 GB> [ - ]
Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM3' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM3'
Successful Full VM restore of Virtual Machine 'VM1'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM1' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM4' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring 1,073,741,824 'VM2' e:\Hyper-V\VM2\Virtual Hard Disks\1G.vhdx
[Done]
Restoring VM configuration information for 'VM4'
Restoring virtual disks for virtual machine 'VM3'
  Total disks to restore       : 1
  Disks to restore concurrently : 1
  Maximum Sessions per disk    : 2

Using 2 session(s) to restore virtual machine 'VM3' Disk 1 .
Restoring 'VM3' Hard Disk 1
  Type           : IFFULL
  Timestamp      : 10/05/2018 14:59:27
  Bytes to restore : 1,073,741,824

Restoring virtual disks for virtual machine 'VM4'
  Total disks to restore       : 1
  Disks to restore concurrently : 1
  Maximum Sessions per disk    : 2

Restoring 0 VM2 [Done]
Using 2 session(s) to restore virtual machine 'VM4' Disk 1 .
Restoring 'VM4' Hard Disk 1
  Type           : IFFULL
  Timestamp      : 10/05/2018 14:59:32
  Bytes to restore : 1,073,741,824

Successful Full VM restore of Virtual Machine 'VM2'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM2' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM5' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM5'
Restoring 1,073,741,824 'VM3' e:\Hyper-V\VM3\Virtual Hard Disks\1G.vhdx

```

```

[Done]
Restoring virtual disks for virtual machine 'VM5'
  Total disks to restore      : 1
  Disks to restore concurrently : 1
  Maximum Sessions per disk   : 2

Using 2 session(s) to restore virtual machine 'VM5' Disk 1 .
Restoring 'VM5' Hard Disk 1
  Type      : IFFULL
  Timestamp  : 10/05/2018 14:59:51
  Bytes to restore : 1,073,741,824

Restoring 1,073,741,824 'VM4' e:\Hyper-V\VM4\Virtual Hard Disks\1G.vhdx
[Done]
Restoring          0    VM3 [Done]
Restoring          0    VM4 [Done]
Successful Full VM restore of Virtual Machine 'VM3'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM3' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Successful Full VM restore of Virtual Machine 'VM4'

Restore processing finished.
Restoring 1,073,741,824 'VM5' e:\Hyper-V\VM5\Virtual Hard Disks\1G.vhdx
[Done]

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM4' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring          0    VM5 [Done]
Successful Full VM restore of Virtual Machine 'VM5'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM5' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Total number of objects restored:      5
Total number of objects failed:        0
Total number of bytes transferred:    5.00 GB
Data transfer time:                    66.95 sec
Network data transfer rate:            78,307.54 KB/sec
Aggregate data transfer rate:          72,000.19 KB/sec
Elapsed processing time:                00:01:12

```

Task

Set a maximum of 2 simultaneous VM restore operations for VMs VM1 and VM2:

```

dsmc restore vm vm=VM1,VM2 -VMMAXRESTOREPARALLELVMs=2
-VMMAXRESTORESessions=10 -VMMAXRESTOREPARALLELDisks=1
-asnode=kingston23_HV_TGT -optfile=dsm.KINGSTON23_HV_DM.opt
-replace=yes

```

This command assigns two simultaneous VM restore operations with 4 sessions per disk per VM, assuming that the VM contains only one disk.

Tip: If you set -VMMAXRESTORESessions=50, there will still be only 4 sessions per disk, as there is only one disk in the VM.

Command output:

```

IBM Storage Protect
Command Line Backup-Archive Client Interface
  Client Version 8, Release 1, Level 7.4
  Client date/time: 10/05/2018 15:21:36
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights
Reserved.

Node Name: KINGSTON23_HV_DM
Session established with server ASANA2: Linux/x86_64
Server Version 7, Release 1, Level 9.000

```

```

Server date/time: 10/05/2018 15:19:44  Last access: 10/05/2018 15:00:10

Accessing as node: KINGSTON23_HV_TGT
Restore function invoked.

Restore VM command started.  Total number of virtual machines to process: 2

Restore of Virtual Machine 'VM1' started

Restore of Virtual Machine 'VM2' started

Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM1' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM1'
Restoring virtual disks for virtual machine 'VM1'
    Total disks to restore      : 1
    Disks to restore concurrently : 1
    Maximum Sessions per disk   : 4

Using 4 session(s) to restore virtual machine 'VM1' Disk 1 .
Restoring 'VM1' Hard Disk 1
    Type      : IFINCR
    Timestamp : 10/05/2018 14:59:23
    Bytes to restore : 1,073,741,824

< 909.25 MB> [ - ]
Starting Full VM restore of 'Hyper-V' Virtual Machine 'VM2' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring VM configuration information for 'VM2'
Restoring 1,073,741,824 'VM1' e:\Hyper-V\VM1\Virtual Machines\1G.vhdx
[Done]
Restoring virtual disks for virtual machine 'VM2'
    Total disks to restore      : 1
    Disks to restore concurrently : 1
    Maximum Sessions per disk   : 4

Using 4 session(s) to restore virtual machine 'VM2' Disk 1 .
Restoring 'VM2' Hard Disk 1
    Type      : IFFULL
    Timestamp : 10/05/2018 14:59:24
    Bytes to restore : 1,073,741,824

Restoring          0    VM1 [Done]
Successful Full VM restore of Virtual Machine 'VM1'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM1' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Restoring 1,073,741,824 'VM2' e:\Hyper-V\VM2\Virtual Hard Disks\1G.vhdx
[Done]
Restoring          0    VM2 [Done]
Successful Full VM restore of Virtual Machine 'VM2'

Restore processing finished.

Successful Full VM restore of 'Hyper-V' Virtual Machine 'VM2' target node
name='KINGSTON23_HV_TGT', data mover node name='KINGSTON23_HV_DM'

Total number of objects restored:      2
Total number of objects failed:      0
Total number of bytes transferred:    2.00 GB
Data transfer time:                   30.08 sec
Network data transfer rate:           69,703.69 KB/sec
Aggregate data transfer rate:         45,697.11 KB/sec
Elapsed processing time:               00:00:45

```

Related tasks

[“Restoring multiple virtual machines in parallel \(optimized restore\)” on page 229](#)

With parallel restore processing, you can restore multiple virtual machines (VMs) at the same time to help optimize the performance of VM restore operations.

Related reference

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Vmmaxrestoresessions” on page 202](#)

The `vmmaxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

[“Vmmaxrestoreparalleldisks” on page 196](#)

Specify the `vmmaxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

[“Optimizing restore operations” on page 227](#)

To help optimize virtual machine (VM) restore operations, adjust the settings for the `vmmaxrestoresessions`, `vmmaxrestoreparallelvms`, and `vmmaxrestoreparalleldisks` options. With these options, you can optimize restore operations while balancing the resources on the Hyper-V hosts.

Vmmaxrestoresessions

The `vmmaxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

An optimized restore operation is one in which parallel restore capability is enabled at the subdisk level of a virtual disk.

At least one session must be allocated for each disk that is being restored.

If the value of `vmmaxrestoresessions` is less than the value of `vmmaxrestoreparalleldisks` multiplied by `vmmaxrestoreparallelvms`, the value will automatically be adjusted to the value of `vmmaxrestoreparalleldisks` multiplied by `vmmaxrestoreparallelvms` at runtime.


Supported clients

This option can be used with supported Windows clients.

Options file

This option is valid in the data mover options file (`dsm.hostname_HV_DM.opt`) or on the command line with the **restore vm** command. It can also be included in a client option set on the IBM Storage Protect server.

Syntax



Parameters

integer

Specifies the number of IBM Storage Protect server sessions that are created during the restore operation. The default is 8. The maximum is 100.

Examples

Options file

VMMAXRESTORES 10

Command line:

dsmc restore vm -VMMAXRESTORES=10

Related tasks

[“Restoring multiple virtual machines in parallel \(optimized restore\)” on page 229](#)

With parallel restore processing, you can restore multiple virtual machines (VMs) at the same time to help optimize the performance of VM restore operations.

Related reference

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Vmmaxrestoreparalleldisks” on page 196](#)

Specify the `vmmaxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

[“Vmmaxrestoreparallelvms” on page 197](#)

The `vmmaxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

[“Optimizing restore operations” on page 227](#)

To help optimize virtual machine (VM) restore operations, adjust the settings for the `vmmaxrestoresessions`, `vmmaxrestoreparallelvms`, and `vmmaxrestoreparalleldisks` options. With these options, you can optimize restore operations while balancing the resources on the Hyper-V hosts.

Vmmaxsnapshotretry

Use the `vmmaxsnapshotretry` option to specify the maximum number of times to retry a snapshot operation of a virtual machine (VM) if the initial snapshot fails with a recoverable condition.

During a VM backup, if a snapshot of a VM fails due to a temporary condition, Data Protection for Microsoft Hyper-V automatically retries the snapshot operation up to the number of times that is specified by the `vmmaxsnapshotretry` option. If the snapshot still fails after the maximum number of retries is reached, the snapshot operation for the VM is not retried and the backup attempt fails.

For example, a recoverable condition might be caused by two backup requests that started at about the same time, backing up VMs that reside on the same volume. One backup operation reports that the snapshot failed because the backup cannot be started while another backup is running for the same VM. In this case, Data Protection for Microsoft Hyper-V will retry the snapshot operation after the first VM backup is completed.

If the initial error is not recoverable, a snapshot is not attempted. For example, if an error occurs with the Volume Shadow Copy Services (VSS) writer during the initial snapshot process, the backup processing stops and Data Protection for Microsoft Hyper-V does not retry the snapshot operation.

This option is valid only for Hyper-V backup operations on Windows Server 2012 and 2012 R2 operating systems.

Supported clients

This option is valid for all supported Windows clients. This option can also be defined on the server.

Options file

This option is valid in the client options file (`dsm.opt`) or on the command line for the **Backup VM** command. It can also be included on the server in a client option set. It cannot be set in the Preferences Editor.

Syntax

➔ `VMMAXSNAshotretry` 

Parameters

integer

Specifies the maximum number of times to retry the snapshot operation of a VM if the initial snapshot attempt fails with a recoverable condition. The default is 20. The maximum is 30. The minimum is 1.

For example, if the `vmmaxsnapshotretry` option is set to 12, Data Protection for Microsoft Hyper-V retries the snapshot operation up to 12 times after the initial snapshot failed during a VM backup operation. If the snapshot still fails after 12 retries are reached, no more retries are attempted, and the backup attempt fails.

At least 10 minutes must elapse before the next snapshot retry attempt. The time between attempts will be longer when the failed VM is part of a snapshot with VMs that are currently being backed up. The backup operation of the other VMs must be completed and the snapshot is removed by the backup operation before a retry attempt can be made.

Examples

Options file

```
vmmaxsna 12
```

Command line

```
dsmc backup vm -vmmaxsna=12
```

Related concepts

[“Tuning scheduled VM backups for Windows Server 2012 and 2012 R2 clusters” on page 59](#)

Beginning with Data Protection for Microsoft Hyper-V version 8.1.2, you can back up more virtual machines (VMs) in parallel and across nodes in a cluster. A cluster node backup operation always retries the snapshot on volumes with snapshots that failed with a recoverable condition. You can also tune the number of VMs in a snapshot to reduce the workload of a snapshot for the Hyper-V host.

Related reference

[“Vmmaxpersnapshot” on page 195](#)

Use the `vmmaxpersnapshot` option to specify the maximum number of virtual machines (VMs) to include in a Hyper-V snapshot. The VMs in the snapshot are backed up to the IBM Storage Protect server.

Vmmaxvirtualdisks

The `vmmaxvirtualdisks` option specifies the maximum size of Hyper-V virtual machine disks (VHDX) to include in a backup operation.

Use the `vmmaxvirtualdisks` option with the `vmskipmaxvirtualdisks` option to specify how the data mover processes large virtual machine (VM) disks during a backup operation:

- Set the `vmmaxvirtualdisks` option to specify the maximum size of the VM disks to include.
- Set the `vmskipmaxvirtualdisks` option to back up the VM disks that do not exceed the maximum size (and exclude any VM disks that exceed the size), or fail the operation.


Supported clients

This option is valid for supported Windows clients that operate as data movers that back up Hyper-V virtual machines.

Options file

Set the `vmmaxvirtualdisks` option in the client options file (`dsm.opt`). You can also specify this option as a command-line parameter on the **backup vm** command.

Syntax



Parameters

size

Specifies the maximum size, in terabytes (TB), of the VM disks to include in a backup operation. The range is an integer 2 - 8; the default is 2. The maximum is 8 TB (equivalent to 8192 GB).

To ensure that the VM disk size that is included in backup operations is always the maximum size, specify 999. Use this value as the most effective method to ensure that the maximum value is always set. This value prevents the need to continuously modify the option files.

When you also specify the `vmskipmaxvirtualdisks yes` option, VM disks that are the specified maximum size or smaller are backed up and VM disks that are larger than the specified maximum size are excluded.

When you also specify the `vmskipmaxvirtualdisks no` option, backup operations fail if a VM disk is larger than the specified maximum size.

Examples

Options file:

```
vmmaxvirtualdisks 3
```

Command line:

Back up VM disks that are 5 TB or smaller and exclude VM disks that are larger than 5 TB:

```
backup vm VM1 -vmmaxvirtualdisks=5 -vmskipmaxvirtualdisks=yes
```

Back up VM disks that are 3 TB or smaller and fail the backup operation if a VM disk is larger than 3 TB:

```
backup vm VM1 -vmmaxvirtualdisks=3 -vmskipmaxvirtualdisks=no
```

Back up VM disks that are 8 TB or smaller and exclude VM disks that are larger than 8 TB:

```
backup vm VM1 -vmmaxvirtualdisks=8 -vmskipmaxvirtualdisks=yes
```

Or:

```
backup vm VM1 -vmmaxvirtualdisks=999 -vmskipmaxvirtualdisks=yes
```

Vmmc

Use the `vmmc` option to store virtual machine backups by using a management class other than the default management class.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax

➤ VMMC — *management_class_name* ➤

Parameters

management_class_name

Specifies a management class that applies to the backed up virtual machine data. If you do not set this option, the default management class of the node is used.

Examples

Task:

Run a backup of the virtual machine that is named `myVirtualMachine` and save the backup according to the management class that is named `myManagmentClass`.

```
dsmc backup vm "myVirtualMachine" -vmmc=myManagmentClass
```

Vmprocessvmwithphysdisks

Use the `vmprocessvmwithphysdisks` option to control whether Hyper-V RCT virtual machine (VM) backups are processed if the VM has one or more physical disks (pass-through disks) provisioned.

A VM can access the storage on a physical disk that is connected directly to the Hyper-V server. This physical disk is called a *pass-through disk*.

When you set this option to `yes`, the data on any physical disks is excluded from backup operations, but the configuration information for the physical disks is saved with the VM backup. During a restore operation, you can restore the physical disk configuration by setting the `vmskipphysdisks no` option. If the original physical disks are available, they are reconnected to the restored VM.

This option is valid only for RCT backups on Windows Server 2016. This option does not apply to Hyper-V VSS backups on Windows Server 2012 or Windows Server 2012 R2.

Supported Clients

This option is valid for clients on Windows Server 2016 or later operating systems.

Options File

Place this option in the client options file (`dsm.opt`) or specify it as a command-line parameter on the **backup vm** command.

Syntax

➤ VMPROCESSVMWITHPHYSDisks  ➤

Parameters

No

The backup operation of the VM fails if one or more physical disks are detected. This value is the default.

Yes

VMs that contain one or more physical disks are backed up. This option backs up the physical disk configuration without backing up the data on the physical disks.

Examples

Options file:

```
VMPROCESSVMWITHPHYSDISKS Yes
```

Command line:

```
dsmc backup vm vmlocal -vmprocessvmwithphysd=yes
```

Related reference

[“Vmskipphysdisks” on page 208](#)

Use the `vmskipphysdisks` option to restore configuration information for physical disks (pass-through disks) that are associated with a Hyper-V virtual machine (VM), if the logical unit numbers (LUNs) that are associated with the volumes on the physical disks are available.

Vmskipmaxvirtualdisks

The `vmskipmaxvirtualdisks` option specifies how backup operations process virtual machine (VM) disks that exceed the maximum disk size.

Use the `vmskipmaxvirtualdisks` option with the `vmmaxvirtualdisks` option to specify how the data mover processes large VM disks during a backup operation:

- Set the `vmskipmaxvirtualdisks` option to back up the VM disks that do not exceed the maximum size (and exclude any VM disks that exceed the size), or fail the operation.
- Set the `vmmaxvirtualdisks` option to specify the maximum size of the VM disks to include.

Supported clients

This option is valid for all supported Windows clients that operate as data movers that back up Hyper-V virtual machines.

Options file

Set the `vmskipmaxvirtualdisks` option in the client options file (`dsm.opt`). You can also specify this option as a command-line parameter on the **backup vm** command.

Syntax



Parameters

No

Specifies that backup operations fail if a virtual machine has one or more VM disks that are larger than the maximum size. This setting is the default value.

Yes

Specifies that backup operations include VM disks that are the maximum size (or smaller) and exclude any VM disks that are larger than the maximum size.

Examples

Options file:

```
vmskipmaxvirtualdisks yes
```

Command line:

Fail a backup operation if a VM disk is larger than 2 TB:

```
backup vm VM1 -vmskipmaxvirtualdisks=no
```

Fail a backup operation if a VM disk is larger than 5 TB:

```
backup vm VM1 -vmskipmaxvirtualdisks=no -vmmaxvirtualdisks=5
```

Back up VM disks that are 8 TB or smaller and exclude VM disks that are larger than 8 TB:

```
backup vm VM1 -vmskipvirtualdisks=yes -vmmaxvirtualdisks=8
```

Vmskipphysdisks

Use the `vmskipphysdisks` option to restore configuration information for physical disks (pass-through disks) that are associated with a Hyper-V virtual machine (VM), if the logical unit numbers (LUNs) that are associated with the volumes on the physical disks are available.

Because physical disks are not included in a VM snapshot, only the configuration information can be restored, and not the data on the volumes.

This option is valid only for restoring Hyper-V VMs on Windows Server 2016. This option does not apply to Hyper-V hosts on Windows Server 2012 or Windows Server 2012 R2.


Supported Clients

This option is valid for clients on Windows Server 2016 or later operating systems.

Options File

Place this option in the client options file (`dsm.opt`), or specify it as a command-line parameter on the **restore vm** command.

Syntax



Parameters

NO

If the original physical disks are available, specify this value to restore the physical disk configuration information that was backed up with the `vmprocessvmwithphysdisks yes` option. The original physical disks are reconnected to the restored VM. If the original physical disks cannot be located, the restore operation fails. This value is the default.

YES

Specify this value if you must restore a VM that you backed up with the `vmprocessvmwithphysdisks yes` option, and the original physical disks cannot be located. This

setting causes the client to skip attempts to locate the physical disks, and does not restore the physical disk configuration information.

Examples

Options file:

```
VMSKIPPHYSDISKS YES
```

Command line:

```
dsmc restore vm vm123 -vmskipphysd=yes
```

Related reference

[“Vmprocessvmwithphysdisks” on page 206](#)

Use the `vmprocessvmwithphysdisks` option to control whether Hyper-V RCT virtual machine (VM) backups are processed if the VM has one or more physical disks (pass-through disks) provisioned.

Vmrestoretype

Use the `vmrestoretype` option with the **query VM** or **restore VM** commands to specify the type of restore operation to perform or query.


Supported Clients

This option can be used with supported Windows clients.

Options file

This option must be specified on the command line of a **restore vm** or **query vm** command. You cannot set this option in the client options file.

Syntax



Parameters

noninstant

Specifies that a regular full VM restore operation is performed. This is the default restore type.

instantaccess

Specifies that a VM backup is restored temporarily. Use this restore type when you want to restore a VM temporarily, to test the integrity of a backup, before you run an instant restore. Any changes that are made to the temporary VM are not saved.

When this restore type is specified on a **query vm** command, the command returns a list of VMs that are running an instant access operation.

vmcleanup

Specifies that a cleanup of the selected VM and its components is performed.

For instant access operations, this option removes the temporary VM and all of its components.

Examples

Command line:

Perform an instant access of the VM named Oslo. The original VM still exists. As a result, the -vmname option is used to assign the new name Oslo_verify.

```
dsmc restore vm Oslo -vmrest=instantaccess -vmname=Oslo_verify
```

Perform a regular (full VM) restore of the virtual machine named San_Jose.

```
dsmc restore vm San_Jose
```

Alternatively, you can also use the following command to run a regular full VM restore operation:

```
dsmc restore vm San_Jose -vmrest=noninstant
```

Run a query to find all active instant access sessions. This query displays all temporary VMs that are running in instant access mode.

```
dsmc query vm * -VMRESToretype=INSTANTAccess
```

Clean up the VM named Oslo and all its components. These components include iSCSI mounts, devices, and temporary data that are associated with the VM name.

```
dsmc restore vm Oslo -vmrest=vmcleanup -vmname=Oslo_verify
```

Related tasks

[“Running instant access operations” on page 125](#)

You can run instant access operations to verify the integrity of a virtual machine (VM) backup or to clean up resources that are created by instant access operations.

Chapter 11. Mount and file restore

The IBM Storage Protect recovery agent 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 for file restore operations.

Tip: Previous versions of Data Protection for Microsoft Hyper-V delivered the function for mounting and snapshot and restoring a file by using 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 Chapter 5, “Getting started with file restore operations,” on page 81.

The content in this “Mount and file restore” collection is provided as a reference for users who prefer the recovery agent method.

IBM Storage Protect recovery agent configurations

The IBM Storage Protect recovery agent provides a variety of configurations for performing file restore and disk / block device exposure.

Off-host file restore

These configurations do not require the IBM Storage Protect recovery agent to be installed in each virtual machine guest. Instead, an off-host instance is responsible for file restore of multiple virtual machines. With this configuration, the mount process exposes a virtual volume from a selected disk partition. For GPT disks, the whole disk must be exposed to make the partitions available, and the disk must be iSCSI connected. Use the recovery agent GUI to accomplish this task.

You must register a node that is associated with the recovery agent. The recovery agent node must be granted proxy authority to access the data node (or nodes) where the snapshots are stored. When a snapshot is mounted to the off-host server, the virtual volume can be network-shared to make it accessible to the virtual machine guest. Or, you can copy the files from the mounted volume to the virtual machine guest by any file-sharing method.

- For step by step restore instructions, see [“Restoring one or more files” on page 214](#)

In-guest file restore

These configurations require IBM Storage Protect recovery agent to be installed in each virtual machine guest. The mount and restore process is performed for a single partition from the backed up disk.

The IBM Storage Protect recovery agent node name is typically granted access only to the virtual machine where it is running with the IBM Storage Protect backup-archive client **dsmtc set access** command. The restore process is typically begun by a user who logs in to the guest machine of the virtual machine.

For these configurations, be sure to compare the specific virtual machine guest operating system requirements with the supported levels of IBM Storage Protect recovery agent. If a specific operating system is not supported, determine if the off-host disk / block device exposure configuration can also be used for file restore. Use the IBM Storage Protect recovery agent GUI to accomplish this task.

- For planning information and operating system-based guidelines, see [Chapter 11, “Mount and file restore,” on page 211](#).
- For step-by-step restore instructions, see [“Restoring one or more files” on page 214](#).

Off-host iSCSI target

This configuration exposes an iSCSI target from the instance of the off-host IBM Storage Protect recovery agent and manually uses an in-guest iSCSI initiator to access the disk snapshot. This configuration requires an iSCSI initiator to be installed within the virtual machine guest. This approach exposes an iSCSI

LUN, rather than the off-host file restore, which exposes an individual disk partition. Use the IBM Storage Protect recovery agent GUI to accomplish this task.

In this configuration, the user specifies the virtual machine guest iSCSI initiator name for the system where the iSCSI device is accessed. After a disk snapshot is mounted, it can be discovered and logged in to by using the iSCSI initiator in the virtual machine guest.

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

1. Mount the GPT disk as an iSCSI target.
 2. Use the Microsoft iSCSI Initiator to log onto the target.
 3. Open the Windows Disk Management to find the disk and bring it online. You can then view the volume in the GPT disk.
- For planning information and operating system-based guidelines, see [Chapter 11, “Mount and file restore,” on page 211](#).
 - For step by step restore instructions, see [“Restoring one or more files” on page 214](#).

Snapshot mount overview

You can use the IBM Storage Protect recovery agent to mount a snapshot and use the snapshot to complete data recovery.

Mount snapshots with the IBM Storage Protect recovery agent GUI. 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 situations when running mount operations:

- When the IBM Storage Protect recovery agent is installed on a guest machine, you cannot start 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 a mount operation. These operations are not allowed because the locking mechanism is for a full virtual machine.
- When you browse the snapshot backup inventory, the operating system version of the virtual machine is the version that was specified when the virtual machine 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.

Mount guidelines

Snapshots can be mounted in either read-only or read/write mode. In read/write mode, 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.

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, 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.

File restore overview

Use the IBM Storage Protect recovery agent for efficient file restore operations and to minimize downtime by mounting snapshots to virtual volumes.

The IBM Storage Protect recovery agent can be used for the following tasks:

- Recovering lost or damaged files from a backup
- Mounting a virtual machine guest volume and creating an archive of the virtual machine guest files
- Mounting database applications for batch reports

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 IBM Storage Protect recovery agent can use a large amount of RAM when it is working in read/write mode.

You can copy the changed files to another volume before you unmount 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 IBM Storage Protect recovery agent mounts snapshots from the IBM Storage Protect server. In the IBM Storage Protect recovery agent GUI, click **Remove** to close an existing connection to the IBM Storage Protect 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 mount machines. You cannot remove the connection to a server when you are running a file restore from that server. You must first dismount all virtual devices and stop all 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 you uninstall the IBM Storage Protect recovery agent. Otherwise, these mounted virtual volumes cannot be unmounted after the IBM Storage Protect 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, consider moving the data to disk or file storage first. 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 NODEDATA** command to move the data back to disk or file storage.

Mounting a snapshot from the same tape storage pool by two instances of Mount can cause one of these results:

- The second Mount instance is blocked until the first instance is complete.
- Both mounts succeed, but the performance is poor.

When restoring data from a mirrored volume, mount only one of the disks that contains the mirrored volume. Mounting both disks causes Windows to attempt a resynchronization of the disks. However, both disks contain a different time stamp 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 from snapshots stored in that node.

File restore guidelines

You can use the IBM Storage Protect recovery agent for efficient file restore and to minimize downtime by mounting snapshots to virtual volumes. File restore is supported from snapshots of NTFS, FAT, or FAT32 volumes.

The mount function cannot be used to mount a snapshot of partitions from a dynamic or GPT-based disk as a virtual volume. Only partitions from an MBR-based, basic disk can be mounted as virtual volumes. File restore from GPT, dynamic, or any other non-MBR or non-basic disk is possible by creating a virtual iSCSI target and using an iSCSI initiator to connect it to your system.

If you are running a file restore of data on dynamic disks, the snapshot must be mounted to a server that has the same version of Windows, or a newer version of Windows, as the node that created the snapshot. Files on the dynamic disk can be accessed indirectly by nodes that have older versions of Windows, by mapping a drive on the older nodes to a CIFS share where the snapshot is mounted.

Important: The ACL values associated with the folders and files that are restored in a file restore operation are not transferred to the restored files. To maintain ACL values, use the **XCOPY** command when copying files from the target.

Restoring one or more files

You can restore one (or more) files from a virtual machine that was backed up to IBM Storage Protect server storage.

Before you begin

If your restore operation accesses the virtual machine disk snapshot with an in-guest iSCSI initiator, make sure the following conditions exist before proceeding:

- The iSCSI device is configured and the iSCSI Initiator program is running.
- Port 3260 is open in the LAN firewall between the system where the IBM Storage Protect recovery agent GUI is installed and the initiator system.

About this task

To mount a backed up virtual machine disk and export the mounted volume for a file restore operation, complete the following steps:

Procedure

1. Start the IBM Storage Protect recovery agent GUI.
On the Windows system, go to **Start > Apps by name > IBM Storage Protect > IBM Storage Protect Recovery Agent**.

The IBM Storage Protect recovery agent GUI can either be installed on the virtual machine guest or installed on a separate host.

2. Connect to the IBM Storage Protect server by clicking **Select IBM Storage Protect server**.
The target node is where the backups are located. You can manage the level of access to the target node data by specifying a different node name in the **Node access method** section.
3. Select a virtual machine from the list.

Tip: You can find your virtual machine quickly by typing the first few letters of the machine name in the edit portion of the list box. The list shows only those machines that match the letters you entered. Machine names are case-sensitive.

A virtual machine might display in the list, but if you select it, the snapshots list might be empty. This situation occurs because of one of the following reasons:

- No snapshots completed successfully for that virtual machine.
- The **Fromnode** option was used and the specified node is not authorized to restore the selected virtual machine.

4. Mount the snapshot through an iSCSI connection:
 - a) Click **Mount** in the IBM Storage Protect recovery agent GUI.
 - b) In the **Select mount destination** dialog, click **Mount as an iSCSI target**.
 - c) Enter the name of the target. This name must be unique for each mount.
 - d) Enter the iSCSI initiator name.

The iSCSI initiator name is shown in the **Configuration** tab in the **iSCSI Initiator Properties** dialog. For example:

```
iqn.1991-05.com.microsoft:hostname
```

5. Complete these steps on the target system where the iSCSI initiator is installed:
 - a) Click the **Targets** tab.
 - b) In the **Quick Connect** section, enter the IP address or host name of the system where the IBM Storage Protect recovery agent GUI is installed.
 - c) Click **Quick Connect**.
 - d) In the **Quick Connect** dialog, select the IP address or host name in the **Discovered targets** field and click **Connect**.
 - e) After **Status - Connected** is shown, click **Done**.
 - f) Go to **Control Panel > Administrative Tools > Computer Management > Storage > Disk Management**.
 - i) If the mounted iSCSI target is listed as **Type=Foreign**, right-click **Foreign Disk** and select **Import Foreign Disks**. The **Foreign Disk Group** is selected. Click **OK**.
 - ii) The next screen shows the type, condition, and size of the Foreign Disk. Click **OK** and wait for the disk to be imported.
 - iii) When the disk import completes, press **F5** (refresh). The mounted iSCSI snapshot is visible and contains an assigned drive letter. If drive letters are not automatically assigned, right-click

the required partition and select **Change Drive Letters or Paths**. Click **Add** and select a drive letter.

6. Select the preferred snapshot date. A list of virtual machine disks that are backed up in the selected snapshot displays. Select a disk and click **Mount**.
7. In the Select Mount Destination dialog, check **Create virtual volume from selected partition**. A list of partitions available on the selected disk is shown. For each partition, its size, label, and file system type are displayed.
 - If the disk is not MBR-based, an error message is displayed.
 - By default, only partitions that can be used for file restore are displayed.
 - To display all partitions that existed on the original disk, clear the **Show only mountable partitions** check box.
8. Select the required partition.

Partitions formatted using unsupported file systems cannot be selected.
9. Specify a drive letter or an empty folder as a mount point for the virtual volume.
10. Click **OK** to create a Virtual Volume that can be used to recover the files.
11. When the Virtual Volume is created, use Windows Explorer to copy the files to your preferred location.

Tip: The ACL values associated with the folders and files that are restored in a file restore operation are not transferred to the restored files. To maintain ACL values, use the **XCOPY** command when copying files from the target.

Related tasks

[“Configuring the IBM Storage Protect recovery agent GUI ” on page 50](#)

You must set up the IBM Storage Protect recovery agent GUI for mount and file restore operations.

[“Manually configuring an iSCSI device” on page 57](#)

You must configure the Windows system that is used during an iSCSI mount operation. The snapshot is mounted from IBM Storage Protect server storage.

Chapter 12. IBM Storage Protect recovery agent commands

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 IBM Storage Protect recovery agent take effect immediately.

You can use the recovery agent CLI to manage only one system running the IBM Storage Protect recovery agent.

On a Windows system, click **Start > Apps by name > IBM Storage Protect > Recovery Agent CLI**.

Mount

Use the **mount** command to complete various IBM Storage Protect 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 IBM Storage Protect recovery agent is running.

Syntax for mounting a disk

```
►► RecoveryAgentShell.exe -c — mount — add — -rep — "tsm: — ip — = ————— IP —————→
                                     |
                                     |———— host_name —————|

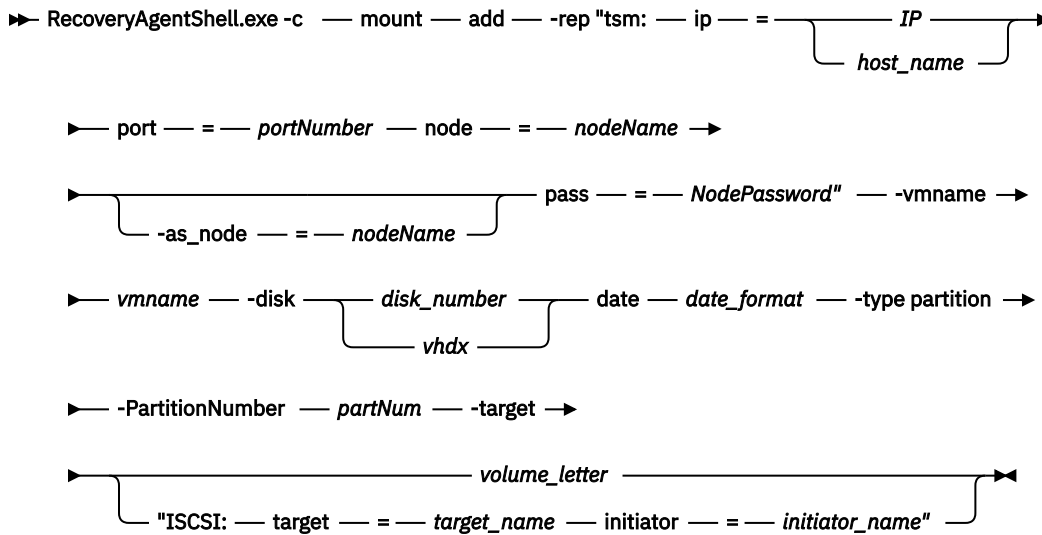
► port — = — portNumber — node — = — nodeName —→

► ————— pass — = — NodePassword" — -vmname —→
|
|———— -as_node — = — nodeName —————|

► vmname — -type — disk — -disk — disk_number — date — date_format — -target —→

► "ISCSI: — target — = — target_name — initiator — = — initiator_name" —►
```

Syntax for mounting partition



Command types

add

Use this command type to mount a disk or volume of a snapshot to the system where IBM Storage Protect recovery agent is running.

The following list identifies the tags and parameters for the **add** command type:

-target

This tag is required. Use this tag to specify the following targets:

- Virtual volume - only for a partition mount
- Reparse point - only for a partition mount
- iSCSI target

-rep

This tag is required. Use it to specify the IBM Storage Protect server that is storing the snapshots, and the IBM Storage Protect node that has access to the 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 IBM Storage Protect 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 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 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 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 IBM Storage Protect 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 IBM Storage Protect 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.

dump

Use this command type to get a list of all the available backups to mount.

The following list identifies the tags and parameters for the **dump** command type:

-rep

This tag is required. Use this tag to specify the IBM Storage Protect server storing the snapshots, and to specify the IBM Storage Protect node that has access to the 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.

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.

The following list identifies the tag for the **remove** command type:

-rep - This tag is required. Use this tag to specify the IBM Storage Protect server connection to be removed.

view

Use this type to view a list of all mounted snapshots. This type has no tags.

Example commands

The following examples use the **-target** tag:

- 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\\"
```

- In the following example an iSCSI target is specified:

```
-target "ISCSI: target=<target_name> initiator=<initiator_name>"
```

In this example, a snapshot of virtual machine named VM-03ent is located on the IBM Storage Protect server with IP 10.10.10.01. Disk number 1 of this snapshot is mounted to the system where the IBM Storage Protect recovery agent is running. The following command 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"
```

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 virtual machine.

```
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>]
```

In the following example, remove the connection to the IBM Storage Protect server (10.10.10.01) using node *NodeName*:

```
mount remove -rep "tsm: NodeName@ip"
```

The following example uses the **view** type:

```
mount view
```

Related links for mounting a Hyper-V snapshot

- [“Set_connection” on page 221](#)
- [“Help” on page 221](#)

Set_connection

The **set_connection** command sets the Recovery Agent CLI to work with a specified IBM Storage Protect recovery agent.

Syntax

➤ RecoveryAgentShell.exe -c — set_connection — mount_computer — *IP address or host_name* ➤

Command type

mount_computer

Use this command type to set the connection from the recovery agent CLI to the system where the IBM Storage Protect recovery agent is installed.

The following list identifies the parameters for the **mount_computer** command type:

IP address or host_name

This variable is required. Specify the IP address or hostname of the system where the IBM Storage Protect recovery agent is installed.

Example commands

In the following example, the recovery agent CLI is set to work with the IBM Storage Protect recovery agent on the *ComputerName* host.

```
set_connection mount_computer ComputerName
```

Related links for setting a connection

- [“Mount” on page 217](#)
- [“Help” on page 221](#)

Help

The **help** command displays the help for all of the supported recovery agent CLI commands.

Syntax

➤ RecoveryAgentShell.exe -c — -h — *command* ➤

Command tag

-h

Use this command tag to show help information.

The following list identifies the parameter for the **mount_computer** command type:

command

This variable is required. Specify the Recovery Agent command for which you want help information.

Example commands

In the following example, the recovery agent CLI is set to work with the IBM Storage Protect recovery agent on the *ComputerName* host.

```
set_connection mount_computer ComputerName
```

Related links for setting a connection

- [“Mount” on page 217](#)
- [“Set_connection” on page 221](#)

Recovery agent command-line interface return codes

Return codes help identify the results of the recovery agent CLI operations.

Use these return codes to check the status of your recovery agent CLI operations.

Table 19. Recovery Agent CLI return codes		
Return Code	Value	Description
0	FBC_MSG_MOUNT_SUCCESS	Command submitted successfully to Data Protection for Microsoft Hyper-V 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.
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 Microsoft Hyper-V 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.

Table 19. Recovery Agent CLI return codes (continued)

Return Code	Value	Description
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 Microsoft Hyper-V 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 Microsoft Hyper-V mount is currently initializing. Try again later.
44	FBC_MSG_CANNOT_LOCK_SNAPSHOT	The snapshot cannot be protected against expiration during this operation. Refer to documentation for more details.

Chapter 13. Optimizing performance

Client options are used with commands to control Data Protection for Microsoft Hyper-V operations. Learn about options and parameters that can help you to tune the performance in your environment.

When you are planning for performance, you typically have goals for Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO). You want to ensure that the performance is good enough to meet your goals, but also balance the resources is used by data protection operations on the Hyper-V hosts in your environment.

The amount of data that is transferred between Data Protection for Microsoft Hyper-V and the IBM Storage Protect server is different for different operations and scenarios:

Backup operations

The amount of data that is transferred for backup operations depends on the state of the backup.

Back up at initial ingestion

When you install Data Protection for Microsoft Hyper-V, no virtual machines (VMs) are backed up and full backup operations are required for each VM. During these initial backup operations, a large amount of data is transferred.

Back up after adding VMs to a Hyper-V host

When you add VMs to a host, the VMs require a full backup. Depending on the number and size of the VMs, a corresponding amount of data is transferred.

In a clustered environment, VMs might move between host nodes. While the total amount of data backed up for the cluster does not change, the amount of data transferred from a particular host changes.

Back up after initial ingestion (steady state)

After the initial full backup operations that occur at initial ingestion, incremental backup operations occur. Incremental backup operations in combination with default-enabled compression and deduplication features result in smaller amounts of data that is transferred. The majority of backup operations are incremental.

Restore operations

Restoring individual or multiple VMs causes a corresponding amount of data to be moved. Consider your RTO and the potential impact to the host and adjust the options.

Disaster recovery operations

In a disaster recovery scenario, you might be more interested in maximizing throughput rather than reducing the impact to the host and might want to adjust the parameters.

Optimizing backup operations

To help optimize virtual machine (VM) backup operations, adjust the settings for the `vmmaxparallel` option. With this option, you can optimize backup operations without adversely affecting Hyper-V hosts.

Overview of the backup options

The following information provides an overview of the `vmmaxparallel` and `vmmaxbackupsessions` options and factors to consider when you are configuring these options for performance.

Tip: Performance testing was conducted to determine the best default value for these options and changes to the values might not be required. For best results, start with the default values and then consider whether your goals are met for backup time as well as for the amount of impact to the Hyper-V hosts.

You might also want to complete backup operations for a limited number of VMs by using the default values and then increase the values if you want to decrease the backup time for the remaining VMs.

In general, increasing the option value results in improved network throughput. However, at a certain point, increasing this value will no longer increase performance. For example, if a 1 GB network is used, when the network is saturated, the throughput will not increase if the option value is increased. In addition, the increased values might result in a higher CPU utilization that negatively affects resource availability on the host.

Also, if network and processor resources are shared with other operations, higher option values might also adversely affect performance on those operations. For example, if the network that is used for backup operations is also used for regular network traffic by the VMs, the VMs might experience slower network speeds during backup operations.

If your priority is to improve throughput, increase the default values. If your priority is to minimize impact to the hosts, decrease the values.

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` depends on the processing power of the Hyper-V host and the performance of I/O between the host and the IBM Storage Protect server.

For example, if you are moving data to the server over a busy or slow LAN, you might have to limit the number of VMs in each parallel backup operation.

Similarly, if the Hyper-V host processing capabilities are limited for any reason, you might also want to limit the number of VMs.

To help optimize performance, use a minimum of a 10GB network in your Data Protection for Microsoft Hyper-V environment.

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 data mover that is running on the Hyper-V host determines the actual number of sessions that are required based on the incoming workload and uses 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. This change ensures that there are as many sessions as there are VMs.

You might have to experiment to find the optimum value for the `vmmaxbackupsessions` option. 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.

The benefit of increasing the `vmmaxbackupsessions` option value is limited if your storage system performance is slower than the available network speed between the data mover and the server.

Restriction: You cannot run multiple backup sessions per VM disk on the Windows Server 2012 operating system. However, you can run multiple backup sessions per VM disk on all other supported Windows operating systems (Windows Server 2012 R2, Windows Server 2016, and later operating systems).

Related tasks

[“Backing up multiple virtual machines in parallel \(optimized backup\)” on page 227](#)

With parallel backup processing, you can back up multiple virtual machines (VMs) at the same time to help optimize the performance of VM backup operations.

Related reference

[“Vmmaxparallel” on page 194](#)

The `vmmaxparallel` option is used to configure parallel backups of several virtual machines (VMs) by using a single instance of the data mover. This option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at a time.

[“Vmmaxbackupsessions” on page 192](#)

The `vmmaxbackupsessions` option specifies the maximum number of IBM Storage Protect server sessions that send virtual machine (VM) data to the server that can be included in an optimized backup operation.

Backing up multiple virtual machines in parallel (optimized backup)

With parallel backup processing, you can back up multiple virtual machines (VMs) at the same time to help optimize the performance of VM backup operations.

Before you begin

Review the information in [“Optimizing backup operations” on page 225](#).

Procedure

Complete the following steps on the data mover system:

1. Go to the `C:\Program Files\IBM\StorageProtect\baclient` directory.
2. Open the data mover options file (`dsm.hostname_HV_DM.opt`) 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
```

4. Issue the **backup vm** command.
For example:

```
dsmc backup vm vm1
```

Using the examples provided, the backup operations for the VM `vm1` at the VM, virtual disk, or subdisk level can include 5 VMs and 10 sessions. Backup operations are limited to 5 VMs per host.

Related reference

[“Vmmaxparallel” on page 194](#)

The `vmmaxparallel` option is used to configure parallel backups of several virtual machines (VMs) by using a single instance of the data mover. This option specifies the maximum number of VMs that can be backed up to the IBM Storage Protect server at a time.

[“Vmmaxbackupsessions” on page 192](#)

The `vmmaxbackupsessions` option specifies the maximum number of IBM Storage Protect server sessions that send virtual machine (VM) data to the server that can be included in an optimized backup operation.

[“Backup VM” on page 147](#)

Use the **backup vm** command to back up Hyper-V virtual machines.

Optimizing restore operations

To help optimize virtual machine (VM) restore operations, adjust the settings for the `vmmaxrestoresessions`, `vmmaxrestoreparallelvms`, and `vmmaxrestoreparalleldisks` options. With these options, you can optimize restore operations while balancing the resources on the Hyper-V hosts.

Overview of the restore options

The following information provides an overview of the `vmmaxrestoresessions`, `vmmaxrestoreparallelvms`, and `vmmaxrestoreparalleldisks` options and factors to consider when you are configuring these options for performance.

Tip: Performance testing was conducted to determine the best default value for these options and changes to the values might not be required. For best results, start with the default values and then consider whether your goals are met for restore time as well as for the amount of impact to the Hyper-V hosts and the IBM Storage Protect server.

You might also want to complete restore operations for a limited number of VMs by using the default values and then increase the values if you want to decrease the restore time for the remaining VMs.

In general, increasing the option value results in improved network throughput. However, at a certain point, increasing this value will no longer increase performance. For example, if a 1 GB network is used, when the network is saturated, the throughput will not increase if the option value is increased. In addition, the increased values might result in a higher CPU utilization that negatively affects resource availability on the host.

If network and processor resources are shared with other operations, higher option values might also adversely affect performance on those operations. For example, if the network that is used for restore operations is also used for regular network traffic by the VMs, the VMs might experience slower network speeds during restore operations.

If your priority is to improve throughput, increase the default values. If your priority is to minimize impact to the hosts, decrease the values.

vmmaxrestoresessions

The `vmmaxrestoresessions` option is used to control the maximum number of IBM Storage Protect server sessions that can be used in an optimized restore operation at one time. Although this option sets the maximum number of sessions that are allowed, the data mover that is running on the Hyper-V host and network speeds determine the actual number of sessions that are required based on the incoming workload.

For maximum performance, the minimum value of the `vmmaxrestoresessions` option must be equal to or greater than the value of the `vmmaxrestoreparalleldisks` option multiplied by the value of the `vmmaxparallelvms` option. If the value of the `vmmaxrestoresessions` option is below this minimum value, it will be increased at runtime to the appropriate value.

You might have to experiment to find the optimum value for the `vmmaxrestoresessions` option. 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 `vmmaxrestoresessions` option.

The benefit of increasing the `vmmaxrestoresessions` option value is limited if your storage system performance is slower than the available network speed between the data mover and the server.

For example, increasing the value of the `vmmaxrestoresessions` option can make better use of the network and disk resources on the systems, though you might have to reduce the number of restore sessions to ensure that the network transfer and disk speeds are not at capacity. If you must use a higher number of restore sessions, consider adding more disks or splitting the VMs across different storage pools or storage devices on the Hyper-V host.

vmmaxrestoreparallelvms

The `vmmaxrestoreparallelvms` option is used to control the maximum number of VMs that can be restored at any one time. The optimal value for `vmmaxrestoreparallelvms` depends on the processing power of the Hyper-V host and the performance of I/O between the host and the IBM Storage Protect server.

For example, if you are moving data from the server over a busy or slow LAN, you might have to limit the number of VMs in each parallel restore operation.

Similarly, if the Hyper-V host processing capabilities are limited for any reason, you might also want to limit the number of VMs.

To help optimize performance, use a minimum of a 10GB network in your Data Protection for Microsoft Hyper-V environment.

vmmxrestoreparalleldisks

The `vmmxrestoreparalleldisks` option is used to control the maximum number of virtual disks that can be restored at the same time per VM, regardless of the value of the `vmmxrestoreparallelvms` option or the number of VMs to be restored. If sufficient sessions are available as specified by the `vmmxrestoresessions` option, the data mover will restore the number of virtual disks per VM at a time as specified by the `vmmxrestoreparalleldisks` option.

Related tasks

[“Restoring multiple virtual machines in parallel \(optimized restore\)” on page 229](#)

With parallel restore processing, you can restore multiple virtual machines (VMs) at the same time to help optimize the performance of VM restore operations.

Related reference

[“Vmmxrestoresessions” on page 202](#)

The `vmmxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

[“Vmmxrestoreparallelvms” on page 197](#)

The `vmmxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

[“Vmmxrestoreparalleldisks” on page 196](#)

Specify the `vmmxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

Restoring multiple virtual machines in parallel (optimized restore)

With parallel restore processing, you can restore multiple virtual machines (VMs) at the same time to help optimize the performance of VM restore operations.

Before you begin

Review the information in [“Optimizing restore operations” on page 227](#).

About this task

Use the following options to optimize the restore operation. Adjust the numbers, depending on the resources available to your environment:

- Use the `vmmxrestoreparallelvms` option to adjust the number of VMs that are restored in parallel.
- Use the `vmmxrestoresessions` option to allocate the number of IBM Storage Protect server sessions for the restore operation and to each VM that is being restored.
- Use the `vmmxrestoreparalleldisks` option to adjust the maximum number of virtual disks to restore in parallel for each VM.

Procedure

Complete the following steps on the data mover command line:

1. Go to the `C:\Program Files\IBM\StorageProtect\baclient` directory.
2. Open the data mover options file (`dsm.hostname_HV_DM.opt`) with your preferred text editor.
3. Enter the option name and one or more blank spaces, followed by the option value.

For example:

```
vmmxrestoreparallelvms 2  
vmmxrestoresessions 8  
vmmxrestoreparalleldisks 2
```

4. Issue the **restore vm** command.

For example:

```
dsmc restore vm vm1,vm2,vm3 -vmname="*_<timestamp>"
```

Using the provided examples, the operation can restore two VMs in parallel, with four maximum sessions available to each VM (eight total restore sessions divided by two VMs), and include a maximum of two disks in parallel per VM. The VMs are restored to new VMs with names that are composed of the original VM name, each appended with the date and time of the restore operation.

Related reference

[“Restore VM” on page 158](#)

Use the **restore vm** command to restore a Microsoft Hyper-V virtual machine (VM) that was previously backed up by IBM Storage Protect for Virtual Environments: Data Protection for Microsoft Hyper-V.

[“Vmmaxrestoreparalleldisks” on page 196](#)

Specify the `vmmaxrestoreparalleldisks` option to restore multiple virtual disks at the same time on a single virtual machine (VM).

[“Vmmaxrestoreparallelvms” on page 197](#)

The `vmmaxrestoreparallelvms` option controls the number of virtual machines (VMs) that can be restored at the same time.

[“Vmmaxrestoresessions” on page 202](#)

The `vmmaxrestoresessions` option defines the maximum number of IBM Storage Protect server sessions that can be included in an optimized restore operation for a virtual machine (VM).

Appendix A. Troubleshooting

Solutions to Data Protection for Microsoft Hyper-V issues are provided.

The following topics are available:

- [“Locating log files” on page 231](#)
- [“Troubleshooting with PowerShell cmdlets” on page 231](#)
- [“Virtual machine backup fails with the 0x800705B4 error in the Hyper-V event log” on page 231](#)
- [“Unsupported characters in virtual machine and Hyper-V host or cluster names” on page 231](#)
- [“The file restore interface shows the wrong drive letter assignments and the system reserved disk” on page 232](#)
- [“An SSL connection cannot be made” on page 232](#)
- [“The SSL certificate for the agent is not valid” on page 232](#)
- [“A VM backup or restore operation cannot start when another VM operation is in progress” on page 233](#)

Locating log files

For information about Data Protection for Microsoft Hyper-V log files, see the following topics:

- [“Data Protection for Microsoft Hyper-V log activity options” on page 50](#)
- [“Trace options for Data Protection for Microsoft Hyper-V” on page 235](#)

Troubleshooting with PowerShell cmdlets

You can troubleshoot Data Protection for Microsoft Hyper-V operations with PowerShell cmdlets. For more information, see [“Troubleshooting Data Protection for Microsoft Hyper-V operations” on page 234](#).

Virtual machine backup fails with the 0x800705B4 error in the Hyper-V event log

During VM backup operations on Windows Server 2016, this error can occur if you run a resilient change tracking (RCT) full backup of a virtual machine (VM) with many VM disks. The snapshot operation either times out or runs out of space on the file space on the server.

If the VM backup operation fails, search the Hyper-V event log for the 0x800705B4 error. If this error is present, complete the following steps to help improve the performance of the snapshot operation:

1. Ensure that the Hyper-V VM is a generation 2 VM.
2. Ensure that only SCSI disks are attached to the generation 2 VM (instead of a mix of SCSI and IDE disks).
3. Move the Hyper-V snapshot folder from the default location (C:\ProgramData\Microsoft\Windows\Hyper-V\Snapshots) to a faster drive that is not the Windows system drive (for example, the D: drive).

Unsupported characters in virtual machine and Hyper-V host or cluster names

Data Protection for Microsoft Hyper-V does not support backing up virtual machines and Hyper-V hosts or clusters 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

The file restore interface shows the wrong drive letter assignments and the system reserved disk

Ensure that the automount feature on Windows is not enabled.

By default, the Data Protection for Microsoft Hyper-V installer automatically disables the automount feature with the **diskpart** command. This action is required to show correct drive letter assignments and to hide the system reserved disk in the IBM Storage Protect file restore interface.

The automount feature was most likely enabled after the installation of Data Protection for Microsoft Hyper-V. Use the **diskpart** command to disable the automount feature.

An SSL connection cannot be made

The following message might appear in the Data Protection for Microsoft Hyper-V Management Console if the SSL certificate is invalid in any way, such as if you reinstalled Data Protection for Microsoft Hyper-V and the old SSL certificate was not deleted.

```
GVM6065E 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-trustore.jks RC=215
```

Delete all the files in the

C:\IBM\StorageProtect\webserver\usr\servers\veProfile\tsmVmGUI\truststores folder. Then, and restart the Data Protection for Microsoft Hyper-V Management Console and run the configuration wizard. Accept the security certificate when prompted.

The SSL certificate for the agent is not valid

You might receive an SSL connection error if the security certificate for the remote client agent is not valid or not up-to-date.

For example, if the certificate files (`dsmcert.sth`, `dsmcert.idx`, and `dsmcert.kdb`) in the `C:\Program Files\Tivoli\TSM\baclient` directory were deleted or are corrupted, the following message appears in the data mover error log (`dsmerror.hostname_HV_DM.log`):

```
ANS1592E Failed to initialize SSL protocol.
```

The method that you use to resolve this problem depends on the level of the IBM Storage Protect server that you are connecting to:

- If you are connecting to an IBM Storage Protect 8.1.2 or later server, or a version 7.1.8 or later V7 server, complete one of the following steps:
 - Stop the client acceptor service on the data mover node and mount proxy node (if file restore is enabled) and re-run the Data Protection for Microsoft Hyper-V configuration wizard on the stand-alone host or on any host in a cluster.

For more information, see [“Configuring Data Protection for Microsoft Hyper-V with the wizard” on page 39](#).
 - Update the node definition on the IBM Storage Protect server by specifying the `SESSIONSECURITY=TRANSITIONAL` parameter. The security certificate is re-created when you sign on to the IBM Storage Protect server from the Data Protection for Microsoft Hyper-V Management Console.

For more information, see [UPDATE NODE](#).
- If you are connecting to an IBM Storage Protect 8.1.1 or earlier V8 server, or a version 7.1.7 or earlier server, see [Dsmcutil commands: Required options and examples](#).

A VM backup or restore operation cannot start when another VM operation is in progress

The following message is displayed if a backup or restore operation is started while another VM operation is in progress:

```
ANS5176W The requested virtual machine operation cannot be performed because a virtual machine backup or restore operation is already in progress. Please retry the operation after the first operation completes.
```

This message appears in the following situations:

- You started a backup or restore operation of a VM and another backup or restore operation is already in progress on the same host.
- You started a backup or restore operation of a VM and another scheduled backup of any VM on the same host is running, or if someone else interactively started the operation from another location.

If you encounter this message, wait for the running operation to finish, then restart your backup or restore operation.

Related reference

[“Troubleshooting application protection of guest virtual machines” on page 121](#)

If Data Protection for Microsoft Hyper-V is configured for application protection of virtual machines (VMs) that host application data and you encounter a problem during VM backup operations, try to reproduce the problem in your environment.

Troubleshooting Data Protection for Microsoft Hyper-V operations

You can retrieve diagnostic information to resolve Data Protection for Microsoft Hyper-V issues by running Microsoft Windows PowerShell cmdlet commands.

Before you begin

Ensure that you prepare your environment to use PowerShell cmdlets. For more information, see [“Preparing to use PowerShell cmdlets with Data Protection for Microsoft Hyper-V” on page 133](#).

Procedure

Complete the following steps on the system where Data Protection for Microsoft Hyper-V is installed.

1. Display log file information in a PowerShell Viewer by issuing the following command:

```
PS C:\> Show-DpHvApiLogEntries
```

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.

2. Display the trace information from a trace file by issuing the following command:

```
PS C:\> Show-DpHvApiTraceEntries
```

3. To gather logs to review detailed diagnostic information parameter or to send to IBM Support, save the logs in a compressed file by issuing the following command:

```
PS C:\> Get-DpHvProblemDeterminationInfo -review
```

By default, this command saves the `DpHvProblemDetermination.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.

4. Optional: Each Data Protection for Microsoft Hyper-V cmdlet provides parameters. To view parameters, issue the following **help** command:

```
help cmdlet name -ShowWindow
```

Related reference

[“Data Protection for Microsoft Hyper-V log activity options” on page 50](#)

The `FRLog.config` options control the content and format of logging information for Data Protection for Microsoft Hyper-V Management Console and file restore operations.

[“Trace options for Data Protection for Microsoft Hyper-V” on page 235](#)

By setting tracing options in the `FRLog.config` file, you can troubleshoot problems that you might encounter during Data Protection for Microsoft Hyper-V and file restore operations.

Trace options for Data Protection for Microsoft Hyper-V

By setting tracing options in the `FRLog.config` file, you can troubleshoot problems that you might encounter during Data Protection for Microsoft Hyper-V and 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\frGUI
```

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/StorageProtect/webserver/usr/servers/veProfile/logs`.

Appendix B. Data Protection for Microsoft Hyper-V messages

Explanations and suggested actions are provided for messages that are issued by Data Protection for Microsoft Hyper-V.

Messages that begin with the GVM prefix are provided in ascending numerical order. In some messages, the explanation and user action are provided in the message itself.

Some messages that begin with the GVM prefix are also shared with IBM Storage Protect for Virtual Environments: Data Protection for VMware.

For messages that begin with the ANS prefix, see [ANS 0000-9999 messages](#).

GVM5900E	The operation failed with return code <i>return code</i>
GVM5901E	An internal error occurred: <i>type of error</i>
GVM5902E	A connection with the IBM Storage Protect server could not be established.

Explanation

The server might not be running.

User response

Check the network connection with the server machine. Verify that the server is running and try to log in again.

GVM5903W	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.

User response

Click OK to delete the data or click Cancel to cancel this action.

GVM5904W	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.

User 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.

GVM5905W	The VM <i>VM name</i> exists, are you going to over-write it?
GVM5906W	The VM <i>VM name</i> is running, make sure the system is powered down, then hit OK to continue.
GVM5907I	A server connection with the name <i>server name</i> has been successfully created. Click OK to continue.
GVM5908W	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.

User 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.

GVM5909I	The VM <i>VM name</i> is spanned into multiple datastores. It can only be restored to its original location.
-----------------	---

GVM5910E	An error occurred while writing to the server's database file, tsmserver.props
-----------------	---

Explanation

The server definition could not be written to the tsmserver.props file.

User response

The file must reside in the install directory of the Data Protection for Virtual Environments. Before you try the action again, verify that the file exists and that the file is not write protected.

GVM5911E	A connection with the vCenter server could not be established.
-----------------	---

Explanation

The server might not be running.

User response

This might indicate a network problem. Ensure that the server is running and the machine is accessible. Try the action again.

GVM5912I	A connection with the vCenter server has been established.
-----------------	---

GVM5913E	The VMCLI inquire configuration command failed, the following messages describe the error.
-----------------	---

Explanation

The Derby database might not be running.

User response

Correct the problem. Try the action again.

GVM5914I	The VMCLI inquire configuration command completed successfully.
-----------------	--

GVM5915E	Failed to determine which product or products are installed.
-----------------	---

Explanation

See message.

User response

Correct the problem. Try the action again.

GVM5916I	Successfully determined which product or products are installed.
-----------------	---

GVM5917E	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.

User response

Select the restore points from the same datacenter or select just a single restore point.

GVM5918E	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.

User response

For restores from IBM Storage Snapshot, all restore points must come from the same backup. You cannot restore multiple VMs that come from of different backups.

GVM5919E	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.

User 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.

GVM5920E	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.

User response

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM5921E	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.

User response

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM5922E	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.

User response

Make sure the correct URL tag is specified with a valid value for the specified mode. Retry accessing the GUI.

GVM5923E	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.

User response

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM5924E	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.

User response

Make sure the tag is specified with a valid value. Retry accessing the GUI.

GVM5925E	The VM name entered <i>VM name</i> conflicts with an existing VM. Please enter a different name.
-----------------	---

GVM5926E	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</i> <i>exception message</i>
-----------------	--

GVM5927E	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.
-----------------	---

GVM5928E	An error occurred while processing the response from the Web server. Detail: <i>error</i>
-----------------	--

GVM5929E	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>
-----------------	---

GVM5930E	No matching device class found. Please return to source page and reselect.
-----------------	---

GVM5931E	No matching proxy node found. Please return to source page and reselect.
-----------------	---

GVM5932E	No proxy ESX hosts available.
-----------------	--------------------------------------

GVM5933I	Password set successfully.
-----------------	-----------------------------------

GVM5934E	Set password failed. Error: <i>message</i>
-----------------	---

Explanation

The password may be incorrect or the server is not running.

User 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.

GVM5935E	Get managed domain failed. Error: <i>message</i>
GVM5936E	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 Snapshot repository.

User response

Select the same type of restore points or select just a single restore point.

GVM5937E	Backup ID is null.
-----------------	---------------------------

Explanation

An internal error occurred.

User response

Refresh the table and perform the action again.

GVM5938E	Task ID is null.
-----------------	-------------------------

Explanation

An internal error occurred.

User response

Refresh the table and perform the action again.

GVM5939E	Could not open a pop-up window.
-----------------	--

Explanation

An internal error occurred.

User response

Try the action again.

GVM5940E	Virtual machine name is null.
-----------------	--------------------------------------

Explanation

An internal error occurred.

User response

Refresh the table and perform the action again.

GVM5941E	Datastore does not exist.
-----------------	----------------------------------

Explanation

An internal error occurred.

User response

Refresh the table and perform the action again.

GVM5942I	No selection was made, the whole virtual machine will be attached.
-----------------	---

Explanation

No selection was made.

User response

Continue with the action or cancel the action.

GVM5943I	Domain set successfully.
GVM5944E	Set domain failed. Error: <i>message</i>

Explanation

The server might not be running.

The permissions on the file directory may be incorrect.

User 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.

GVM5945E	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>
-----------------	---

GVM5946E	The schedule requires use of the following datacenters that are not known to the system.
-----------------	---

Datacenters: *list*
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: *summary*

GVM5947E The schedule requires use of the following hosts that are not known to the system.
Hosts: *list*
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: *summary*

GVM5948E The schedule requires use of the following datastores that are not known to the system.
Datastores: *list*
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: *summary*

GVM5949E The schedule requires use of the following virtual machines that are not known to the system.
Virtual Machines: *list*
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*

GVM5950I Password set successfully.
Warning: *message*

Explanation

The password was set successfully with a warning.

User response

Follow the action described in the warning message.

GVM5951E 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: *error*

GVM5952E The following command requires confirmation from the server:
"*Command*"

Explanation

A command was issued, and a reply was expected. Some commands require a confirmation, which you cannot issue through the Data Protection for Virtual Environments GUI.

User response

Issue the command from the command line.

GVM5953E The following command is unknown to the server:
"*Command*"

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.

User response

Verify that the command is valid for the server version and platform, and verify that the command syntax is correct.

GVM5954E The syntax of the following command is incorrect:
"*Command*".

Explanation

See message.

User 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.

GVM5955E An internal server error occurred.

Explanation

See message.

User 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.

GVM5956E	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.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5957E	The database recovery log is full.
-----------------	---

Explanation

See message.

User 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.

GVM5958E	The server database is full.
-----------------	-------------------------------------

Explanation

See message.

User response

Before trying the action again, extend the server database. Contact the administrator of the IBM Storage Protect server.

GVM5959E	The server is out of storage space.
-----------------	--

Explanation

See message.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5960E	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.
-----------------	---

GVM5961E	The object that you are attempting to access does not exist on the server.
-----------------	---

GVM5962E	The object that you are attempting to access is currently in use by another session or process. Retry the action at a later time.
-----------------	--

GVM5963E	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.
-----------------	--

GVM5964E	The object that you are attempting to access or remove is not available.
-----------------	---

Explanation

See message.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5965E	The server encountered an I/O error while processing the request. For more information, see the operating system event or error log.
-----------------	---

GVM5966E	The action failed because the transaction could not be committed.
-----------------	--

Explanation

See message.

User response

Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5967E	The action failed because of a resource lock conflict.
-----------------	---

Explanation

See message.

User response

Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5968E	The action failed because of a mode conflict.
-----------------	--

Explanation

See message.

User response

Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5969E	The action failed because the server could not start a new thread.
-----------------	---

Explanation

See message.

User response

Retry the action at a later time. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5970E	The server is not licensed to perform this action. If a license was purchased, use the command line to register the license.
-----------------	---

GVM5971E	The specified destination is not valid.
-----------------	--

Explanation

See message.

User response

Enter a different destination or update the configuration with a valid destination, and try the action again.

GVM5972E	The specified input file cannot be opened. Verify the file name and directory permissions, then try the action again.
-----------------	--

GVM5973E	The specified output file cannot be opened. Verify the file name and directory permissions, then try the action again.
-----------------	---

GVM5974E	An error occurred while writing to the specified output file.
-----------------	--

Explanation

See message.

User response

Check the file system to ensure that there is enough space. Check the operating system event or error log for more information.

GVM5975E	The specified administrator is not defined to this server.
-----------------	---

Explanation

See message.

User response

Ensure that the administrator name was entered correctly. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5976E	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.

User response

Correct the SQL query and try again.

GVM5977E	This operation is not allowed with this object.
-----------------	--

Explanation

See message.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5978E	The table was not found in the server database.
-----------------	--

Explanation

See message.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5979E	The specified file space name is not compatible with the filespace type.
-----------------	---

Explanation

Unicode file space names are incompatible with non-unicode names.

User response

Enter a file space name of the correct type and try the action again.

GVM5980E	The specified TCP/IP address is not valid. Verify the TCP/IP address and try the action again.
GVM5981E	No objects were found that match the search conditions.
GVM5982E	Your administrative ID on this server is locked. An administrator with system authority can unlock your ID.
GVM5983E	The connection to the server was lost while performing the action.

Explanation

See message.

User response

This might indicate a network problem. Ensure that the server is running and the machine is accessible. Retry the action.

GVM5984E	Your ID or password is not valid for this server.
-----------------	--

Explanation

See message.

User response

Enter a valid ID or password for your IBM Storage Protect Server.

GVM5985E	Your password expired on this server.
-----------------	--

Explanation

Your IBM Storage Protect password has expired.

User response

Reset your password on the IBM Storage Protect Server or contact your IBM Storage Protect Server administrator to reset it.

GVM5986E	The server cannot accept new sessions. If sessions are disabled for this server, issue the ENABLE SESSIONS command from the command line.
-----------------	---

GVM5987E	A communications failure occurred while processing the request. Retry the action at a later time.
-----------------	--

GVM5988E	The administrative API encountered an internal error while processing the request.
-----------------	---

GVM5989E	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.

User response

Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5990E	The following command contains one or more invalid parameters: <code>""command""</code>.
-----------------	---

Explanation

The Data Protection for Virtual Environments GUI tried to run a command, but the call to the API contained one or more invalid parameters.

User 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.

GVM5991E	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.

User 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.

GVM5992E	The administrator's authority level on this server cannot be determined.
-----------------	---

Explanation

See message.

User response

Use a different administrator ID. Before trying the action again, contact the administrator of the IBM Storage Protect server.

GVM5993E	An object with the name that you specified already exists on the server. Enter a different name.
-----------------	---

GVM5994E	The version of the server is not supported by the Data Protection for Virtual Environments GUI.
-----------------	--

GVM5995E	An internal error has occurred.
-----------------	--

Explanation

The operation failed after encountering an internal error.

User 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.

GVM5996E	The operation failed, please go to the log for more details.
-----------------	---

GVM5997E	Wrong format of the end date and time. Please enter the end date and time format as yyyyMMddHHmmss.
-----------------	--

GVM5998E	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.

GVM5999E	The ESXHOST name you entered is too long. Please change to a shorter one.
-----------------	--

GVM6000E	Wrong Backup ID. Please try again.
-----------------	---

GVM6001E	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.

GVM6002E	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.

GVM6003E	Wrong format of the start date and time. Please enter the start date and time format as yyyyMMddHHmmss.
-----------------	--

GVM6004I	Backup task <i>Task Name</i> started, would you like to monitor this task now?
-----------------	---

GVM6005I	Delete backup task completed successfully.
-----------------	---

GVM6006E	Delete backup task failed, please check log for more detail.
-----------------	---

GVM6007I	Restore Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
-----------------	--

GVM6008E	<i>Error Or Warning</i>
-----------------	--------------------------------

GVM6009I	Mounted backup Item could not be restored.
-----------------	---

GVM6010I	Result of attach is <i>status</i> (Task ID: <i>Task ID</i>), refer to events list to get the details.
-----------------	--

GVM6011I	Result of detach is <i>status</i> (Task ID: <i>Task ID</i>), refer to events list to get the details.
-----------------	--

GVM6012I	Command successfully submitted to the IBM Storage Protect server. Detail: <i>Server Messages</i>
-----------------	---

GVM6013E	The command submitted to the IBM Storage Protect server failed.
-----------------	--

**Error: Error Code
Error Messages**

Explanation

The cause of the problem is identified in the message text.

User response

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM6014E	No IBM Storage Protect server connection, please configure the IBM Storage Protect server in the configuration panel.
GVM6015E	The selected items can only be under ONE datacenter.
GVM6018E	The virtual machine <i>VM name</i> exists. Delete the virtual machine first before restoring it.
GVM6019E	The target virtual machine <i>VM name</i> is running. Close the virtual machine before restoring virtual disks to it.
GVM6020E	Some of selected virtual disks exist in target virtual machine. Remove those virtual disks from target virtual machine before restoring to it.
GVM6021E	A VMCLI command failed. Error: Error Messages

Explanation

The cause of the problem is identified in the message text.

User response

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM6023E	A command submitted to the IBM Storage Protect server failed. Error: Error Messages
-----------------	--

Explanation

The cause of the problem is identified in the message text.

User response

Correct the problem based on the information that is provided in the message text. Then, try the action again.

GVM6024E	Cannot find the file with format 'summary.date.log' in the path: <i>path</i>
GVM6025E	Cannot find the IBM Storage Snapshot installation path using the VMCLI inquire_config command.
GVM6026E	A VMCLI command to get version failed.
GVM6027I	Backup task <i>Task ID</i> started, would you like to monitor this task now?
GVM6028E	The Data Protection for Virtual Environments Web Server could not be contacted.

Explanation

The Data Protection for Virtual Environments GUI has attempted to contact its Web Server. The operation was not successful.

User response

Perform one or more of the following steps to try and determine the problem:

- Verify that the Data Protection for Virtual Environments 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 Data Protection for Virtual Environments GUI. Start the GUI again when the problem is resolved.

GVM6029I	Command successfully submitted to the server.
GVM6030E	No host is found in datacenter <i>datacenter name</i>. Select another datacenter to restore.
GVM6031W	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 the Data Protection for Virtual Environments GUI.

User response

This schedule must be modified outside the the Data Protection for Virtual Environments GUI.

GVM6032W	One or more VMs exist. Do you want to continue the restore operation and overwrite the existing VMs?
GVM6033E	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.

User 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.

GVM6034E	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.

User response

Pick another node name to use. If you want to re-use this node, then unselect the 'Register Node' checkbox.

GVM6035E	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.

User 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.

GVM6036E	The passwords in the entry field and the verify field do not match. Please try again.
-----------------	--

Explanation

The new passwords entered do not match.

User response

Clear the fields and enter the same password in both password fields.

GVM6037W	Please select one or more Datacenters to be managed.
-----------------	---

Explanation

At least one Datacenter must be selected.

User response

Add one or more Datacenter(s) into the Managed Datacenters list.

GVM6038W	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.

User response

Assign a password for nodes that are to be registered.

GVM6039I	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.
-----------------	--

GVM6040I	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.

GVM6041I	This task was skipped because it was not necessary or a pre-requisite task failed.
-----------------	---

GVM6042E	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.

User response

Try the operation again.

GVM6043I	Managed datacenters have changed. Please go to the data mover page to verify or change your current mappings.
-----------------	--

GVM6044I	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.
-----------------	--

GVM6045E	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 rule(s).

User response

Try with another password.

GVM6046W	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.

User response

We strongly recommend you use the Configuration Wizard with a proper IBM Storage Protect Administrative ID.

GVM6047W	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.

User response

Try using another node name.

GVM6048W	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.

User response

Try using another node name.

GVM6049E	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.

User response

Try with another password that only contain valid characters.

GVM6050E	The password entered is not acceptable on this Server because of the reason below. Choose another password. Error: <i>message</i>
-----------------	--

Explanation

IBM Storage Protect Server could not accept the password chosen. The reason why this password is not valid is given in the message.

User response

Try with another password that meets the rule(s).

GVM6051E	Filter has changed, select Apply filter before continuing.
-----------------	---

Explanation

Filter pattern must be applied after it is changed.

User response

Click the Apply filter button.

GVM6052E	Select at least one item from a datacenter to continue.
-----------------	--

Explanation

A host, host cluster, or VM must be selected to do a backup.

User response

Select an item under a datacenter.

GVM6053E	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.

User response

Create multiple backup tasks, with less selected items per task.

GVM6054I	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.

User response

Select OK to proceed, or select Cancel to retain all selections.

GVM6055E	Datacenter node <i>datacenter node name</i> 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.

User 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.

GVM6056E	IBM Storage Protect datacenter node <i>datacenter node name</i> maps to vCenter datacenter name <i>datacenter name</i> in the vmcli configuration file, but <i>datacenter name</i> 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.

User 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.

GVM6057E	You have selected items from multiple datacenters: <i>datacenter list</i>. 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.

User response

Check and correct the selections to make sure all selections are under the same datacenter.

GVM6058E	The selected items <i>item list</i> are not found under datacenter <i>datacenter name</i> 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.

User 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.

GVM6062E	The password entered is not acceptable on this Server because it is too short. Passwords must have a least <i>minPasswordLength</i> characters.
-----------------	--

Explanation

IBM Storage Protect Server could not accept the password chosen because it is too short.

User response

Try with another password that is longer than the required minimum length.

GVM6063E	<i>Component</i> is downlevel, so its use is disabled in the GUI. You will only be able to use the GUI for <i>component</i>.
-----------------	---

GVM6064E	Mismatching IBM Storage Protect Server entries in the current settings is detected. IBM Storage Protect Server definition used by the GUI: <i>server1</i> IBM Storage Protect Server where backups are stored: <i>server2</i> 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 Data Protection for Virtual Environments 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.

User 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.

GVM6065E	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.

User response

Check the TSM-ve-truststore.jks for a valid certificate, ensure TSM-ve-truststore.jks is in the correct default location.

GVM6066E	The password entered is not acceptable on this Server because it is too long. Passwords cannot have more than <i>maxPasswordLength</i> characters.
-----------------	---

Explanation

IBM Storage Protect Server could not accept the password chosen because it is too long.

User response

Try with another password that is shorter than the allowed maximum length.

GVM6067E	The SSL Connection could not be made. The IBM Storage Protect SSL certificate is invalid.
-----------------	--

Explanation

IBM Storage Protect Server did not accept the SSL connection. The TSM-ve-truststore.jks has an invalid IBM Storage Protect SSL certificate.

User 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.

GVM6068E	The non-SSL connection could not be made. This IBM Storage Protect Admin ID requires a 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.

User 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.

GVM6069E	Your selections have caused the backup task definition to require <i>count</i> 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.

User response

De-select the newly added virtual machines checkbox or create multiple backup tasks with less selected items per task.

GVM6070E	Your selection of virtual machines has caused the backup task definition to require <i>count</i> 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.

User 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.)

GVM6071E	There is no data mover node proxy relationship for datacenter node <i>datacenter node name</i>. Review the data mover relationships on the Configuration tab or the IBM Storage Protect server.
-----------------	--

GVM6072E	There is no datacenter node defined for datacenter <i>datacenter name</i>. Review the node configuration on the Configuration tab.
-----------------	---

GVM6073I	Node <i>name name</i> is currently locked. The Configuration Wizard will attempt to unlock this node if you choose to continue.
-----------------	--

GVM6074E	A connection with the IBM Storage Protect server (<i>Address:Port</i>) could not be established. Please verify the server address and admin port <i>Server or Admin port</i> are correct.
-----------------	--

Explanation

The server might not be running or specified admin port or server admin port may be incorrect.

User 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 address and admin port information is correct.

GVM6075E	The vCenter user name or password is not valid. Please try again.
-----------------	--

Explanation

The vCenter user name or password is not valid.

User response

Enter the user name or password again.

GVM6076E	Permission to perform this operation was denied. Please try with other user name.
-----------------	--

Explanation

The vCenter user name is not valid.

User response

Enter another user name.

GVM6077I	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.
-----------------	--

GVM6078W	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: <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.
-----------------	--

GVM6079I	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.
-----------------	--

GVM6080I	ID has been changed without save. Previous ID will be loaded.
-----------------	--

GVM6081I	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.
-----------------	---

GVM6082I	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.
-----------------	--

GVM6083I	There are non-English characters contained in one or more datacenters. The domain will be adjusted accordingly.
-----------------	--

GVM6084E	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.

User response

Datacenter will not be added to the domain.

GVM6085W	Node <i>Node Name</i> already exists on the server. Attempt to rename node to <i>New Node Name</i>?
-----------------	--

Explanation

Node name is already registered on the IBM Storage Protect server.

User response

Click Yes to attempt to rename node. Click No to make other changes. Example: unclick register node, rename node manually.

GVM6086W	The following virtual machines for host <i>Host Name</i> have unsupported characters in their name: <i>Invalid Virtual Machine Names</i>. 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: "" ' ; * ? , < > / |

User response

Rename the identified virtual machines to remove unsupported characters from their name.

GVM6087E	The following host clusters have unsupported characters in their name: <i>Invalid Host Clusters</i>. 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: "" ' ; * ? , < > / |

User response

Rename the identified host clusters to remove unsupported characters from their name. Or, remove them from your backup selection.

GVM6088E	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.

User response

Rename the identified virtual machines to remove unsupported characters from their name. Or, select different virtual machines to back up.

GVM6089E	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: "" ' ; < > / |

User response

Change the filter pattern to remove unsupported characters, then apply the filter again.

GVM6090E	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.

User response

Add a datastore to the destination ESX host. Then, select this datastore as the temporary restore destination.

GVM6091E	There was an error creating opt file: <i>file name</i>.
-----------------	--

Explanation

An error was encountered when trying to write to file.

User response

Try the operation again.

GVM6092E	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.

User response

Check environment and ensure user has proper rights before trying operation again.

GVM6093E	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.

User 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.

GVM6094W	Local services were setup successfully but were unable to verify firewall access for these executable files: <i>agentExe</i> <i>cadExe</i> <i>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.

User 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.

GVM6095E	Data mover node <i>node name</i> was successfully registered on the server, however no services were created.
-----------------	--

Explanation

An error has occurred when trying to create services for specified node.

User response

Check environment and ensure user has proper rights before trying operation again.

GVM6096E	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.

User response

Log into the host machine specified and view the error log for more information.

GVM6097W	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.

User 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.

GVM6098W	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.

User response

Use the Configuration Settings to reset the VMCLI node password.

GVM6099W	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.

User response

Cancel the selected task or allow the task to continue processing.

GVM6100W	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.

User 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.

GVM6101W	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.

User response

Click 'Dismount' to dismount (cleanup) the instant access virtual machines.

GVM6102E	Selecting multiple virtual machines with different restore types is not allowed.
-----------------	---

Explanation

Restoring multiple virtual machines with different restore types is not supported.

User response

Select virtual machines that have the same restore type.

GVM6103I	Cleanup Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
-----------------	--

GVM6104W	Are you sure that you want to cancel this task?
-----------------	--

Explanation

A cancel task command is submitted. Refresh to see the cancel progress.

User response

Cancel the selected task or allow the task to continue processing.

GVM6105I	Your current UI role does not allow you to view backup property notebook.
-----------------	--

GVM6106I	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.
-----------------	---

GVM6107E	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.

User response

Log into the host machine where data mover resides and view the error log for more information.

GVM6108W	Login information for vCenter needed.
-----------------	--

Explanation

In order to install new local dm services, vCenter credentials are needed.

User response

Enter vCenter credentials in order to continue.

GVM6109E	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.

User response

Add the required privileges for the user.

GVM6110E	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.

User response

Add the required permissions for the user.

GVM6111I	A new data center (<i>name</i>) was detected. Go to the Data Mover Nodes page to add a data center node for it.
-----------------	--

GVM6112W	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.

User response

Click 'Dismount' to dismount (cleanup) the mounts and shares.

GVM6113I	Dismount Task <i>Task ID</i> is started successfully, would you like to monitor this task now?
-----------------	---

GVM6114W	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.

User response

Make sure the option file was deleted. If it still exists, delete this file manually.

GVM6115W	The remove operation for IBM Storage Protect service: <i>service</i> failed.
-----------------	---

Explanation

An error prevented the IBM Storage Protect service from being removed.

User response

Check the environment and ensure that the user has sufficient rights to run this operation. Then, try the operation again.

GVM6116E	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.

User response

Start the iSCSI service manually.

GVM6117E	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.

User 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...' checkbox is selected. The server's certificate must be

retrieved and a truststore created using the procedure that is documented in the 'Learn more...' link.

GVM6118E	You have selected organization VDCs from more than one provider VDC. For backup tasks, all selected organization VDCs must belong to the same provider VDC. Change your selections and retry the operation.
-----------------	--

GVM6119E	The following vcloud resources(vApp, organization, organization vDC) are invalid for selection because they have unsupported characters in their name: <i>reslist</i>
-----------------	--

Explanation

In order to create backup tasks, vcloud resources names must not contain any of the following characters: "" ' : ; * ? , < > / | .

User response

Rename the identified resources to remove unsupported characters from their name. Or, remove them from your backup selection.

GVM6120E	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.
-----------------	--

GVM6121E	The vApp <i>vApp name</i> exists. Choose a different vApp name to be the target of the restore.
-----------------	--

GVM6122E	Your selection of items to back up has caused the backup task definition to require <i>count</i> 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.

User response

Create multiple backup tasks with less items per task

GVM6123E	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.
-----------------	--

GVM6124E	The nodename <i>node name</i> 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.

User response

Pick another node name to use. If you want to re-use this existing node, then unselect the 'Register Node' checkbox.

GVM6125W	Are you certain that you want to remove the data mover node <i>node name</i>?
-----------------	--

GVM6126W	The IBM Storage Protect node <i>node</i> has already been used. If you want a different name other than the default name, edit this field again.
-----------------	---

Explanation

The node is already being used in this configuration.

User response

Try using another node name.

GVM6127E	The Organization VDC node can not be registered because its provider VDC is not valid.
-----------------	---

GVM6128E	The Organization VDC name <i>OVDC name</i> is invalid. For information about supported characters, refer to the IBM Storage Protect Administrator's Reference publication section about naming IBM Storage Protect objects.
-----------------	--

GVM6129I	This task was skipped because it was not necessary. No further action is required.
-----------------	---

GVM6130W	Internet explorer version <i>version</i> 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.

User response

Use a supported version of Internet Explorer or another browser. Supported browser versions are documented in the online help.

GVM6131W	The browser <i>version</i> 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.

User response

Use a supported browser. Supported browser versions are documented in the online help.

GVM6132E	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: <i>VM name</i>
-----------------	---

Explanation

When restoring to an alternate location, the destination virtual machine must not already exist.

User response

Use the single virtual machine restore wizard so that you can rename the destination virtual machine.

GVM6133W	Target datastore not found, select a different destination datastore.
GVM6134E	The user <i>User Name</i> is not authorized to any managed datacenters. Contact your system administrator.
GVM6135E	You do not have required permissions to view virtual machines for this Event.
GVM6136E	You do not have required permissions to view restore points for this virtual machine.
GVM6137E	You do not have required permissions to view some attached points.
GVM6138E	You do not have required permissions to view restore points for this datastore.
GVM6139E	You do not have required permissions to detach for the restore point.
GVM6140E	An error occurred processing user permissions. Contact your system administrator.
GVM6141I	Some datacenters are not shown due to permissions requirements.
GVM6142E	You do not have permissions to cancel this task.
GVM6143I	The task is still in the starting state, please refresh the task and try the cancel again.
GVM6147I	Some datacenters are not available because they have the same name for one or more datacenters. Datacenters with the same name are not supported.
GVM6148E	Windows domain credentials are incorrect. Open the Configuration wizard, 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 wizard is incorrect.

System action

Processing stops.

User response

Run the Configuration wizard again and re-enter the correct Windows domain credentials.

GVM6149E	This action cannot be performed because there is not a VMCLI node defined. To resolve, use the configuration wizard to define the VMCLI node and complete the other steps in the wizard.
GVM6150E	This action cannot be performed because there is not a vCloud Director node defined. To resolve, use the configuration wizard to define the vCloud Director node and complete the other steps in the wizard.
GVM6151E	This action cannot be performed because the connection to the IBM Storage Protect Server is not operational. Correct the connection problem, and retry this action.
GVM6152E	This task requires use of the provider VDC node <i>provider VDC node name</i> from IBM Storage Protect, but this node is not mapped to a known provider VDC in the vCloud Director. This task may not be updated, instead create a new task without dependence on this provider VDC.
GVM6153E	The Organization VDCs listed below were selected but are not configured to the IBM Storage Protect server. You must remove these selections in order to execute this action. <i>org VDC name</i>
GVM6154I	Your current UI role does not allow you to view node details.
GVM6155E	An error occurred when connecting to the IBM Storage Protect server <i>server name</i>. 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.

User response

Launch the Configuration Editor from the Configuration Tab and enter a valid ID or password for your IBM Storage Protect Server.

GVM6156E	The password for the administrative user ID <i>admin id</i> expired on the IBM Storage Protect server <i>server name</i>.
-----------------	--

Explanation

Your IBM Storage Protect administrative password has expired.

User response

Contact your IBM Storage Protect Server administrator to reset the password for the administrative user ID.

GVM6157E	The IBM Storage Protect server port number <i>tcp port</i> is incorrect. The expected value for this port is <i>tcp port from query</i>, which is the value of the TCPPORT option. Please enter the expected value using the configuration wizard.
-----------------	---

Explanation

The value entered in the IBM Storage Protect server port field must match the TCPPORT option on the IBM Storage Protect server.

User response

Use the configuration wizard to change the IBM Storage Protect server port field to the correct value.

GVM6159E	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.
-----------------	---

GVM6160E	An error occurred while writing to the frConfig.props configuration file.
-----------------	--

Explanation

The frConfig.props file contains configuration options for file level restore processing. Possible reasons for this error include the following situations:

- The frConfig.props file is not in the Data Protection for Virtual Environments installation directory.
- The frConfig.props file is write-protected.

System action

Processing stops.

User response

Verify that the file exists in the Data Protection for Virtual Environments installation directory and that the file is not write-protected.

GVM6161E	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.

User response

Disable the file level restore feature. Then, choose whether you want to remove the mount proxy node pair.

GVM6162E	An error occurred while reading the frConfig.props configuration file.
-----------------	---

Explanation

The frConfig.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.

System action

Processing stops.

User response

Verify that the file is not read-protected.

GVM6164W	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.

User response

If this message was not presented as part of using the configuration wizard, it must be retrieved and a truststore created using the procedure that is documented in the help.

GVM6165E	The specified target node '<i>node-name</i>' does not match the node '<i>node-name</i>' stored in the user session.
-----------------	--

Explanation

The input target node to the configuration host operation does not match the target node stored in the connected session.

User response

Retry the operation with the correct target node name.

GVM6166E	A user session is invalid or no SSL certificate to accept.
-----------------	---

Explanation

The initial IBM Storage Protect server connection detects that it requires an SSL certificate and the operation must be called again with the same connection. In this case, the connection is null or invalid.

User response

Make sure the operation call the second time to accept the certificate is using the same initial connection.

GVM6167E	A Windows mount proxy node and a Linux mount proxy node are required to enable File Restore.
-----------------	---

Explanation

Either one mount proxy node or no proxy node were specified for the configure host operation.

User response

Retry the operation with a node list that have a Windows mount proxy node and a Linux mount proxy node.

GVM6168E	Configure host failed. Check the tasks list for more information.
-----------------	--

Explanation

Configuring the host consists of registering the target node, registering the data mover and creating the

services for backup and restore, registering the mount proxy nodes and creating the services for file level restore. One of these tasks encountered an error.

User response

Fix the error and retry the operation.

GVM6169E	Unexpected error while configuring to the IBM Storage Protect server.
-----------------	--

Explanation

Possible reasons for this error include the following situations:

- Unknown error while trying to connecting to the IBM Storage Protect server.
- Unknown error while trying to write to the server's database file, tsmserver.props.

User 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 information is correct.

GVM6170E	Unexpected error, can not get policy domain for node '<i>node-name</i>'.
-----------------	---

Explanation

The target node does not exist on the IBM Storage Protect server or an internal error occurred during the node query.

User response

Run the configuration wizard to register the target node or update the node to another policy domain.

GVM6171E	Unexpected error, schedule '<i>schedule-name</i>' does not exist on the IBM Storage Protect server.
-----------------	--

Explanation

The schedule may have been deleted accidentally during the operation.

User response

Select a different schedule.

GVM6172E	'<i>domain-name</i>' is not a valid Windows domain.
-----------------	--

Explanation

LOCALHOST or the computer name are not valid domains.

User response

Enter a valid domain.

GVM6173E	The domain is missing from the user name.
-----------------	--

Explanation

The user name you entered is not part of a domain.

User response

Ensure that the user name is in the DOMAIN\UserName format.

GVM6174E	The following addresses cannot be reached: <i>httpurl</i>, <i>httpsurl</i>. Verify that the TSM Client Acceptor (CAD) is up and running.
-----------------	---

Explanation

The CAD service is not running for the data mover.

System action

The operation cannot continue without a connection to the data mover CAD service.

User response

Make sure the data mover CAD service is running and that the node has the proper proxy relationships established.

GVM6175E	The TCP port from the HTTP response cannot be retrieved. Verify that the TSM Client Acceptor (CAD) is up and running.
-----------------	--

Explanation

The CAD service is not running for the data mover.

System action

The operation cannot continue without a connection to the data mover CAD service.

User response

Make sure the data mover CAD service is running and that the node has the proper proxy relationships established.

GVM6176E	The TCP port from the HTTP response cannot be parsed or found.
-----------------	---

Explanation

The HTTP stream from the agent does not contain the TCP port number.

System action

The operation cannot continue without a connection to the data mover CAD service.

User response

Make sure the data mover CAD service is running and that the node has the proper proxy relationships established.

GVM6177E	An exception was encountered while parsing TCP port string: <i>tcpport</i>.
-----------------	--

Explanation

The HTTP stream from the agent returned an invalid TCP port number.

System action

The operation cannot continue without a connection to the data mover CAD service.

User response

Make sure the data mover CAD service is running and that the node has the proper proxy relationships established.

Appendix C. 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

A

- accessibility features [263](#)
- actions pane
 - Data Protection for Microsoft Hyper-V Management Console [71](#)
- ad hoc backups
 - Data Protection for Microsoft Hyper-V Management Console [77](#)
- application protection
 - general help [122](#)
 - overview [87](#)
 - troubleshooting [121](#)
 - troubleshooting VSS backup and restore operations [122](#)
- application protection for Exchange Server
 - configuring after VM name change [94](#)
 - install and configure Data Protection for Microsoft Exchange Server [89](#)
- application protection for Microsoft Exchange Server
 - back up data [96](#)
 - configure Data Protection for Microsoft Hyper-V for application protection [90](#)
 - getting started overview [87](#)
 - getting started step 1 [88](#)
 - getting started step 2 [89](#)
 - getting started step 3 [90](#)
 - getting started step 4 [94](#)
 - install and configure Data Protection for Microsoft Hyper-V [88](#)
 - installing and configuring [87](#)
 - overview [87](#)
 - restore a database [94](#)
 - restore backups of other VMs [100](#)
 - restore data [99](#)
 - restore data with cmdlets [102](#)
 - restore data with the command line [100](#)
 - restore databases with the GUI [99](#)
 - restore mailbox data [100](#)
 - scheduling backups [96](#)
 - show file space information [102](#)
 - start iSCSI service [99](#)
 - update mailbox history information [97](#)
 - verify backups [97](#)
 - verify that volumes are not excluded [98](#)
- application protection for Microsoft SQL Server
 - back up data [112](#)
 - Configure Data Protection for Microsoft Hyper-V for application protection [107](#)
 - getting started overview [103](#)
 - getting started step 1 [104](#)
 - getting started step 2 [105](#)
 - getting started step 3 [107](#)
 - getting started step 4 [110](#)
 - install and configure Data Protection for Microsoft Hyper-V [104](#)
 - install and configure Data Protection for Microsoft SQL Server [105](#)

- application protection for Microsoft SQL Server (*continued*)
 - installing and configuring [103](#)
 - manage backup versions [113](#)
 - overview [103](#)
 - restore data [115](#)
 - restore databases with the command line [116](#)
 - restore databases with the GUI [115](#)
 - restore relocated and deleted databases [118](#)
 - restore SQL Server log backups [118](#)
 - restoring a database [110](#)
 - scheduling backups [112](#)
 - scheduling SQL Server log backups [112](#)
 - script for validating VM backups [119](#)
 - show file space information [120](#)
 - start iSCSI service [115](#)
 - verify backups [113](#)
 - verify that volumes are not excluded [114](#)
- application protection for SQL Server
 - configuring after VM name change [110](#)
- automated client failover
 - overview [9](#)

B

- back up
 - multiple VMs [227](#)
 - parallel [192](#), [194](#), [227](#)
- back up data
 - protecting Exchange Server data [96](#)
 - protecting SQL Server data [112](#)
- back up Hyper-V VMs [1](#)
- back up VMs now
 - Data Protection for Microsoft Hyper-V Management Console [77](#)
- backing up
 - large VHDXs [204](#), [207](#)
 - VHDXs up to 8 TB [204](#), [207](#)
- backup
 - incremental forever
 - description [8](#)
 - limitations [10](#)
 - performance, optimizing [225](#)
 - policy management [8](#)
 - RCT backup
 - description [2](#)
 - user interfaces
 - description [4](#)
 - VSS backup
 - description [1](#)
- backup history
 - Data Protection for Microsoft Hyper-V Management Console [75](#)
- backup status
 - Data Protection for Microsoft Hyper-V Management Console [75](#)
- backup vm command [147](#)
- best practices

best practices (*continued*)
excluding VMs [79](#)

C

check configuration
 Linux mount proxy for file restore [47](#)
cluster backups on Windows Server 2012
 reducing schedule contention [59](#)
commands
 backup vm [147](#)
 expire [154](#)
 mount [217](#)
 query VM [154](#)
 restore vm [158](#)
 set_connection [221](#)
communication ports [16](#)
configuration
 advanced tasks [58](#)
configuration wizard
 clusters [39](#)
 file restore [39](#)
 stand-alone host [39](#)
configure
 default port numbers [58](#)
 file restore options [48](#)
 Linux mount proxy for file restore [45](#)
configuring
 clusters [39](#)
 file restore [39](#), [44](#)
 IBM Storage Protect recovery agent GUI [50](#)
 initial configuration [39](#)
 iSCSI mount [57](#)
 overview [39](#)
 security settings [39](#), [42](#)
 stand-alone host [39](#)
configuring after VM name change
 protecting Exchange Server data [94](#)
 protecting SQL Server data [110](#)
configuring TLS
 enable secure communication with the server [54–56](#)
control files [190](#)
customizing
 nodes [20](#)

D

data move node
 overview [6](#)
Data Protection for Microsoft Hyper-V
 comparability [17](#)
 configuring tracing [235](#)
 installing on Server Core [32](#)
 troubleshooting
 diagnostic procedure [234](#)
 uninstalling on Server Core [33](#)
 upgrading [17](#)
 upgrading with the installation wizard [23](#)
 using Data Protection for Microsoft Exchange Server [122](#)
Data Protection for Microsoft Hyper-V
 features
 installable [15](#)

Data Protection for Microsoft Hyper-V Management
 Console

 actions pane [71](#)
 as a snap-in [61](#)
 configuring logging [49](#)
 description [62](#)
 instant access cleanup [131](#)
 instant access operations [130](#)
 logging in [61](#)
 logging options [50](#)
 navigation pane [62](#)
 overview [61](#)
 restoring a VM [78](#)
 results pane [63](#)
 running an ad hoc backup [77](#)
 Schedule History view [66](#)
 set backup policy [72](#)
 setting schedules [72](#)
 setting the at-risk policy [74](#)
 starting [61](#)
 Tasks pane [68](#)
 verify configuration [71](#)
 view backup history [75](#)
 view backup status [75](#)
 viewing schedule history [75](#)
 Virtual Machines view [63](#)

Data Protection for Microsoft Hyper-V Management Console
reconnecting [61](#)

Data Protection for Microsoft Hyper-V overview [1](#)

date format

 specifying [167](#)

dateformat option [167](#)

description

 Data Protection for Microsoft Hyper-V Management

 Console [62](#)

detail option [169](#)

disability [263](#)

disk space requirements

 Windows client [15](#)

documentation [11](#)

domain

 include for full vm backups [169](#)

domain.vmfull option [169](#)

E

enable secure communication with the server
 configuring TLS [54–56](#)

errors [231](#)

exclude

 EXCLUDE.VMDISK [172](#)

EXCLUDE.VMDISK [172](#)

excluding VMs

 best practice [79](#)

expire command [154](#)

F

failover

 client [9](#)

file restore

 check configuration of Linux mount proxy [47](#)

 common tasks [81](#)

file restore *(continued)*

- configure Linux mount proxy [45](#)
- configuring [44](#)
- configuring logging [49](#)
- configuring options [48](#)
- configuring tracing [235](#)
- description [81](#)
- installing Linux mount proxy [34](#)
- installing Linux mount proxy in silent mode [36](#)
- installing Linux mount proxy overview [33](#)
- logging in [84](#)
- logging options [50](#)
- options [48](#)
- prerequisites [82](#)
- procedure [84](#)
- removing [38](#)
- roles [81](#)
- uninstalling Linux mount proxy [37](#)
- upgrading Linux mount proxy [33](#)

file space [169](#)

files

- restore overview [213](#)
- restore task (Windows) [214](#)

G

getting started

- protecting Exchange Server data overview [87](#)
- protecting Exchange Server data step 1 [88](#)
- protecting Exchange Server data step 2 [89](#)
- protecting Exchange Server data step 3 [90](#)
- protecting Exchange Server data step 4 [94](#)
- protecting Microsoft SQL Server data step 1 [104](#)
- protecting SQL Server data overview [103](#)
- protecting SQL Server data step 2 [105](#)
- protecting SQL Server data step 3 [107](#)
- protecting SQL Server data step 4 [110](#)

group backup

- display active and inactive objects [174](#)

H

hardware requirements

- Windows client [15](#)

hyper v [9](#)

Hyper-V [9](#)

Hyper-V cmdlets [8](#)

Hyper-V snapshots

- deleting [8](#)
- rolling back [8](#)

I

IBM Documentation [vii](#)

IBM Storage Protect recovery agent GUI

- configuring [50](#)
- options [50](#)

import security certificate

- for servers earlier than version 8.1.2 or version 7.1.8 [43](#)
- for servers later than version 8.1.2 or version 7.1.8 [39](#)

inactive option [174](#)

include

- INCLUDE.VMDISK [176](#)

include.vm option [174](#)

INCLUDE.VMDISK [176](#)

include.vmsnapshotattempts option [177](#)

include.vmtsmvss option [179](#)

incremental forever

- description [8](#)

installable features

- Data Protection for Microsoft Hyper-V [15](#)

installation package

- download [23](#)

installation procedure

- data mover [28](#)
- Data Protection for Microsoft Hyper-V Management Console [26](#)
- download package [23](#)
- overview [23](#)
- planning [15](#)
- silent [30](#)
- typical [24](#)

installing

- Linux mount proxy for file restore [34](#)
- Linux mount proxy overview [33](#)
- on Server Core systems [32](#)
- planning [15](#)
- security certificate on host [61](#), [133](#)

installing and upgrading

- overview [15](#)

installing in silent mode

- Linux mount proxy for file restore [36](#)

instant access

- cleanup [128](#)
- host restart restriction [132](#)
- remove resources [128](#)
- verify VM backup [125](#)

iSCSI mount

- configuring [57](#)

K

keyboard [263](#)

L

LAN environment [212](#)

limitations on Hyper-V backup operations [10](#)

logging

- Data Protection for Microsoft Hyper-V Management Console [49](#)
- file restore [49](#)

logging in Data Protection for Microsoft Hyper-V Management Console [61](#)

logs

- truncating application logs [179](#)

M

mailbox history information

- updating in Microsoft Exchange Server backups [97](#)

manage backup versions

- protecting SQL Server data [113](#)

management class [8](#)

managing snapshots [8](#)

- maximum VHDX size
 - how to process [207](#)
 - specifying [204](#)
- Mbobjrefreshthresh [183](#)
- Mbpctrefreshthresh [183](#)
- memory requirements
 - Windows client [15](#)
- messages
 - ANS prefix [237](#)
 - Data Protection for Microsoft Hyper-V [237](#)
 - GVM prefix [237](#)
- Microsoft Exchange Server backups
 - updating mailbox history [97](#)
- migrating
 - nodes [18](#)
- mode option [182](#)
- mount and file restore [211](#)
- mount command [217](#)
- mount proxy node
 - overview [6](#)
- mounting snapshots [212](#)

N

- navigation pane
 - Data Protection for Microsoft Hyper-V Management Console [62](#)
- new features in Data Protection for Microsoft Hyper-V version 8.1.27 [ix](#)
- nodes
 - customizing [20](#)
 - migrating [18](#)
 - overview [6](#)
 - prefix [20](#)
 - renaming [18](#)
 - suffix [20](#)
 - updating [18](#)
- noprompt option [184](#)
- numberformat
 - specifying [184](#)
- numberformat option [184](#)

O

- online help
 - PowerShell cmdlets [135](#)
- options
 - dateformat [167](#)
 - detail [169](#)
 - domain.vmfll [169](#)
 - EXCLUDE.VMDISK [172](#)
 - file restore [48](#)
 - inactive [174](#)
 - include.vm [174](#)
 - INCLUDE.VMDISK [176](#)
 - include.vmsnapshotattempts [177](#)
 - include.vmtsmvss [179](#)
 - mbobjrefreshthresh [183](#)
 - mbpctrefreshthresh [183](#)
 - mode [182](#)
 - noprompt [184](#)
 - numberformat [184](#)

- options (*continued*)
 - pick [186](#)
 - pitdate [186](#)
 - pittime [187](#)
 - skipsystemexclude [187](#)
 - timeformat [188](#)
 - vmautostartvm [189](#)
 - vmbackdir [190](#)
 - vmbackupupdateguid [148](#)
 - vmiscsiserveraddress [191](#)
 - vmmaxbackupsessions [192](#)
 - vmmaxparallel [194](#)
 - vmmaxparallelrestorevms [197](#)
 - vmmaxpersnapshot [195](#)
 - vmmaxrestoreparalleldisks [196](#)
 - vmmaxrestoresessions [202](#)
 - vmmaxsnapshotretry [203](#)
 - vmmaxvirtualdisks [204](#)
 - vmmc [206](#)
 - vmprocessvmwithphysdisks [206](#)
 - vmrestoretype [209](#)
 - vmskipmaxvirtualdisks [207](#)
 - vmskipphysdisks [208](#)
- options reference [167](#)
- overview
 - application protection [87](#)
 - back up Hyper-V VMs [1](#)
 - Data Protection for Microsoft Hyper-V [1](#)
 - Data Protection for Microsoft Hyper-V Management Console [61](#)
 - Hyper-V environment [4](#)
 - nodes [6](#)
 - policy management [8](#)
 - protecting Exchange Server data [87](#)
 - protecting SQL Server data [103](#)
 - restore Hyper-V VMs [3](#)
 - user interfaces [4](#)
 - VM backups with RCT [2](#)
 - VM backups with VSS [1](#)

P

- parallel backups [192, 194](#)
- performance, optimizing
 - backup operations [225](#)
 - restore operations [227](#)
- pick option [186](#)
- pitdate [186](#)
- pittime option [187](#)
- port numbers
 - configure [58](#)
- PowerShell cmdlets
 - getting help [135](#)
 - list [135](#)
 - prerequisite steps [133](#)
 - protect VMs [135, 138](#)
 - tasks [138](#)
 - using [133](#)
- prerequisite steps
 - PowerShell cmdlets [133](#)
- preview
 - restore vm [164](#)
- problem determination [231](#)
- protect VMs

protect VMs (*continued*)

PowerShell cmdlets [135](#), [138](#)

publications [vii](#)

Q

query

backups, establish point-in-time [186](#), [187](#)

display active and inactive objects [174](#)

query VM command [154](#)

quiesce applications [179](#)

R

RCT backups

description [2](#)

features [2](#)

migrating to [2](#)

upgrade considerations [22](#)

reconnect to

Data Protection for Microsoft Hyper-V Management Console [61](#)

remove an instant access VM

Data Protection for Microsoft Hyper-V Management Console [131](#)

removing file restore [38](#)

renaming

nodes [18](#), [20](#)

requirements

communication ports [16](#)

resilient change tracking (RCT) backups [2](#)

restore

backups, establish point-in-time [186](#), [187](#)

configuring logging [49](#)

configuring options [48](#)

create list of backup versions to [186](#)

display active and inactive objects [174](#)

file [48–50](#), [82](#), [84](#)

file restore description [81](#)

file restore roles [81](#)

file restore tasks [81](#)

Hyper-V VMs

description [3](#)

logging in [84](#)

multiple VMs [229](#)

options [48](#), [50](#)

parallel VMs [229](#)

performance, optimizing [227](#)

prerequisites [82](#)

procedure [84](#)

user interfaces

description [4](#)

restore a VM

Data Protection for Microsoft Hyper-V Management Console [78](#)

restore backups of other VMs

protecting Exchange Server data [100](#)

restore data

protecting SQL Server data [115](#)

restore data with cmdlets

protecting Exchange Server data [102](#)

restore data with the command line

protecting Exchange Server data [100](#)

restore databases with the command line

protecting SQL Server data [116](#)

restore databases with the GUI

protecting Exchange Server data [99](#)

protecting Exchange SQL data [115](#)

restore mailbox data

protecting Exchange Server data [100](#)

restore relocated and deleted databases

protecting SQL Server data [118](#)

restore SQL Server log backups

protecting SQL Server data [118](#)

restore vm command

preview [164](#)

restoring data

protecting Exchange Server data [99](#)

results pane

Data Protection for Microsoft Hyper-V Management Console [63](#)

running instant access operations

Data Protection for Microsoft Hyper-V Management Console [130](#)

S

Schedule History view

Data Protection for Microsoft Hyper-V Management Console [66](#)

scheduled cluster backups on Windows Server 2012

tuning [59](#)

scheduling backups

protecting Exchange Server data [96](#)

protecting SQL Server data [112](#)

scheduling SQL Server log backups

protecting SQL Server data [112](#)

script for validating VM backups

protecting SQL Server data [119](#)

security certificate on host

installing [61](#), [133](#)

security settings

configuring [42](#)

to connect to servers earlier than version 8.1.2 or version 7.1.8 [43](#)

to connect to servers later than version 8.1.2 or version 7.1.8 [39](#)

self-contained application protection [179](#)

set backup policy

Data Protection for Microsoft Hyper-V Management Console [72](#)

set_connection command [221](#)

setting schedules

Data Protection for Microsoft Hyper-V Management Console [72](#)

setting the at-risk policy

Data Protection for Microsoft Hyper-V Management Console [74](#)

show file space information

protecting Exchange Server data [102](#)

protecting SQL Server data [120](#)

silent install

data mover only [30](#)

Data Protection for Microsoft Hyper-V Management Console only [30](#)

silent uninstall [31](#)

skipssystemexclude [187](#)

- snap-in
 - Data Protection for Microsoft Hyper-V Management Console [61](#)
- snapshot management [8](#)
- snapshots
 - mounting [212](#)
- specify maximum VHDX size [204](#)
- specify the host name or the IP address of the iSCSI server [191](#)
- SSL
 - configuring [54–56](#)
- start iSCSI service
 - protecting Exchange Server data [99](#)
 - protecting SQL Server data [115](#)
- starting
 - Data Protection for Microsoft Hyper-V Management Console [61](#)
- syntax diagram
 - reading [145](#)
 - repeating values [145](#)
 - required choices [145](#)
- system state
 - display active and inactive objects [174](#)

T

- target node
 - overview [6](#)
- Tasks pane
 - Data Protection for Microsoft Hyper-V Management Console [68](#)
- the host name or the IP address of the iSCSI server
 - specifying [191](#)
- time format
 - specifying [188](#)
- timeformat option [188](#)
- tracing
 - configuring [235](#)
 - options [235](#)
- troubleshooting
 - application protection [121](#)
 - VSS backup and restore operations [122](#)

U

- uninstallation procedure
 - silent [31](#)
- uninstalling
 - Linux mount proxy [37](#)
 - on Server Core systems [33](#)
- update mailbox history information
 - protecting Exchange Server data [97](#)
- updating
 - nodes [18](#)
- upgrade by using the installation wizard [23](#)
- upgrade tasks [17](#)
- upgrading
 - Linux mount proxy for file restore [33](#)
 - RCT backups [22](#)
 - version compatibility [17](#)
- use PowerShell cmdlets [133](#)
- using the GUI [61](#)

V

- verify backups
 - protecting Exchange Server data [97](#)
 - protecting SQL Server data [113](#)
- verify configuration
 - Data Protection for Microsoft Hyper-V Management Console [71](#)
- verify that volumes are not excluded
 - protecting Exchange Server data [98](#)
 - protecting SQL Server data [114](#)
- view schedule history
 - Data Protection for Microsoft Hyper-V Management Console [75](#)
- Virtual Machines view
 - Data Protection for Microsoft Hyper-V Management Console [63](#)
- vmauto startvm [189](#)
- vmbackdir option [190](#)
- vmbackupupdateguid option [148](#)
- vmctlmc option
 - options
 - vmctlmc [190](#)
- vmiscsiserveraddress option [191](#)
- vmmaxbackupsessions option [192](#)
- vmmaxparallel option [194](#)
- vmmaxparallelrestorevms option [197](#)
- vmmaxpersnapshot option [195](#)
- vmmaxrestoreparalleldisks option [196](#)
- vmmaxrestoresessions option [202](#)
- vmmaxsnapshotretry [203](#)
- vmmaxvirtualdisks option [204](#)
- vmmc option [8](#), [206](#)
- vmprocessvmwithphysdisks option [206](#)
- vmrestoretype option [209](#)
- vmskipmaxvirtualdisks option [207](#)
- vmskipphysdisks option [208](#)
- Volume Shadow Copy Service (VSS) backups
 - description [1](#)
- volumes
 - restore overview [213](#)
 - restore task (Windows) [214](#)
- VSS backup Data Protection for Microsoft Hyper-V
 - with Data Protection for Microsoft Exchange Server [122](#)
- VSS backups
 - description [1](#)

W

- what's new for version 8.1.27 [ix](#)
- Windows client
 - disk space requirements [15](#)
 - hardware requirements [15](#)
 - memory requirements [15](#)



Product Number: 5725-X00