



Carbonite Server Backup

Hyper-V Agent 9.1

User Guide



© 2023 Open Text. All rights reserved.

One or more patents may cover this product. For more information, please visit <https://www.opentext.com/patents>.

For terms and conditions, see <https://www.carbonite.com/terms-of-use/carbonite-general-enterprise-terms-of-service/>.

Carbonite makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Carbonite reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Carbonite to notify any person of such revision of changes. All companies, names and data used in examples herein are fictitious unless otherwise noted.

No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval System or translated into any language including computer language, in any form or by any means electronic, mechanic, magnetic, optical, chemical or otherwise without prior written permission of:

Carbonite LLC
251 Little Falls Drive
Wilmington, DE 19808
www.carbonite.com

Carbonite and the Carbonite logo are trademarks of Carbonite, LLC. Product names that include the Carbonite mark are trademarks of Carbonite, LLC. All other products or company names mentioned in this document are trademarks or registered trademarks of their respective owners.

Version History

Version	Date	Description
1	July 2022	Initial guide for Hyper-V Agent 9.1.
2	December 2022	In Introduction to the Hyper-V Agent , added note that Hyper-V Agent version 9.12 can back up and restore VMs in Microsoft Azure Stack HCI clusters.
3	July 2023	Restore Hyper-V files, folders and database items now indicates that, to restore files and folders from Windows VMs, the Hyper-V Agent Management service must be installed on the same Windows version or a later version than is installed on the Windows VMs.

Contents

1 Introduction to the Hyper-V Agent	5
1.1 Hyper-V Agent components	5
2 Prepare for a Hyper-V Agent deployment	7
2.1 Portal for managing a Hyper-V Agent	7
2.2 Vaults for Hyper-V backups	7
2.3 Recommended deployment for protecting a Hyper-V standalone host	7
2.4 Recommended deployment for protecting a Hyper-V cluster	8
2.5 Hyper-V Rapid VM Restore requirements	9
2.6 Hyper-V Agent ports	10
2.7 Best practices in a protected Hyper-V environment	11
3 Install, upgrade and uninstall the Hyper-V Agent	13
3.1 Install the Hyper-V Agent Management service	13
3.2 Configure a new Hyper-V Agent	16
3.3 Install the Hyper-V Agent Host service	18
3.4 Upgrade the Hyper-V Agent	21
3.5 Uninstall the Hyper-V Agent Management service	25
3.6 Uninstall the Hyper-V Agent Host service	26
4 Configure a Hyper-V Agent	27
4.1 Change credentials or the network address for accessing Hyper-V	27
4.2 Add vault settings	28
4.3 Add a description	30
4.4 Add retention types	30
4.5 Set up email notifications for backups on a computer	32
4.6 Configure bandwidth throttling	33
4.7 Resolve certificate failures	34
5 Add and schedule a Hyper-V backup job	36
5.1 Best practices for backing up Hyper-V VMs	37
5.2 Best practices for seeding Hyper-V VM backups	37
5.3 Application-consistent backups on Hyper-V VMs	38
5.4 Add a Hyper-V backup job	39
5.5 Edit a Hyper-V backup job	43
5.6 Add or edit a schedule for a Hyper-V backup job	46
5.7 Disable or enable all scheduled backup jobs	49
5.8 Run an ad-hoc backup	50

6 Restore Hyper-V data	52
6.1 Restore Hyper-V VMs	52
6.2 Restore a Hyper-V VM within minutes using Rapid VM Restore	55
6.3 Restore Hyper-V files, folders and database items	61
7 Recover jobs and settings from an offline Hyper-V Agent	66
7.1 Hyper-V disaster recovery and migration	69
8 Delete jobs and computers, and delete data from vaults	73
8.1 Delete a backup job without deleting data from vaults	73
8.2 Delete a backup job and delete job data from vaults	74
8.3 Cancel a scheduled job data deletion	76
8.4 Delete a computer without deleting data from vaults	77
8.5 Undelete Hyper-V environments	78
8.6 Delete a computer and delete computer data from vaults	78
8.7 Cancel a scheduled computer data deletion	80
9 Monitor computers, jobs and processes	82
9.1 Monitor backups and computers using the Current Snapshot	82
9.2 Monitor storage usage using Site Usage charts and emailed alerts	83
9.3 View computer and job status information	85
9.4 View an unconfigured computer's logs	87
9.5 View current process information for a job	88
9.6 View a job's process logs and safeset information	89
9.7 View, export and email backup statuses on the Monitor page	90
9.8 View a Hyper-V VM's backup history and logs	93
9.9 Hyper-V Agent logs and configuration files	95
10 Understanding and troubleshooting Hyper-V processes	96
10.1 Some VMs backed up before others	96
10.2 VM skipped during backup	96
10.3 VM backup fails	96
10.4 Live migration fails	96
10.5 VM restore fails	96
11 Carbonite Server Backup Support	97
11.1 Contacting Carbonite	97
Appendix: Understanding Hyper-V backups on a vault	98
Determine the name of a VM's task on the vault	98

1 Introduction to the Hyper-V Agent

The Hyper-V Agent provides data protection for virtual machines (VMs) in Microsoft Hyper-V environments. The Agent protects VMs in standalone and clustered Hyper-V environments, without requiring agent software to be installed on individual VMs.

Note: Beginning in Hyper-V Agent version 9.12, the Agent can also back up and restore VMs in Microsoft Azure Stack HCI clusters. For more information, see the *Hyper-V Agent release notes*.

The Hyper-V Agent concurrently backs up multiple VMs in a single backup job. In a cluster, backup operations can be distributed across nodes, making the solution scalable in large environments. Within a Hyper-V cluster, the Agent can back up VMs that have migrated to different nodes or to different storage.

You can include multiple VMs in a single backup job, but each VM is backed up as a separate task on the vault. As a result, each VM has a single backup history, even if it is moved from one backup job to another over time. When restoring a VM, you do not need to remember which backup job it was in.

To improve the performance of incremental backups, Hyper-V Agent 9.1 determines which parts of a VM disk have changed since the last backup and only reads disk blocks that have changed. In previous versions, the Agent had to read all disk blocks in an incremental backup.

Note: To determine which disk blocks have changed, the Hyper-V Agent uses Resilient Change Tracking (RCT): a Hyper-V feature that tracks changes on VM disks. Because RCT is only available in Windows Server 2016 or later, the Hyper-V Agent reads all disk blocks when backing up VMs in Hyper-V on Windows Server 2012 R2.

You can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows VMs. Application-consistent backups minimize the amount of work needed to fully restore applications.

You can restore entire VMs using the Hyper-V Agent or restore specific files, folders and database items from Windows VMs. Beginning with Hyper-V Agent 9.00, you can restore a VM within minutes using the Rapid VM Restore method. See [Restore Hyper-V data](#).

1.1 Hyper-V Agent components

As shown in [Components for protecting a Hyper-V environment](#), the Hyper-V Agent consists of two components:

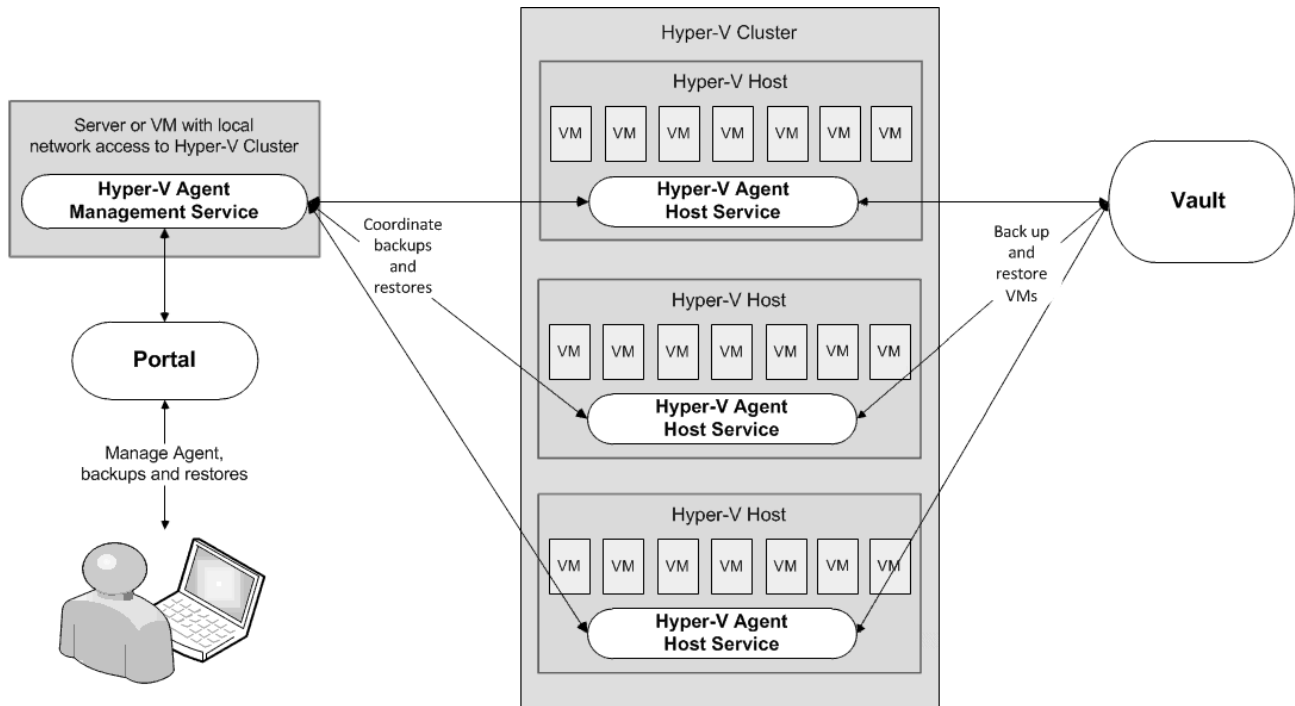
- Hyper-V Agent Management service. The Management service is a central management component that communicates with Portal and delegates backup and restore operations to Hyper-V Agent Host services. The Management service is the only Hyper-V Agent component that directly communicates with Portal.
- Hyper-V Agent Host service. The Host service is installed on one or more hosts in a Hyper-V environment. Host services perform VM backups and restores, as delegated by the Management service. Host services do not directly communicate with Portal so there is no need to open ports

between Host services and Portal. When performing a backup or restore, the Host service communicates directly with the vault where the VM is backed up.

The Hyper-V Agent is closely integrated with Portal. You must use Portal to manage the Hyper-V Agent, back up VMs to a secure vault, and restore VMs. The Portal instance can be hosted by your service provider or installed on-premises.

Even though it consists of more than one component, the Hyper-V Agent appears as a single system in Portal.

Components for protecting a Hyper-V environment



2 Prepare for a Hyper-V Agent deployment

Before installing a Hyper-V Agent, you must do the following:

- Obtain a Portal account for managing the Agent. See [Portal for managing a Hyper-V Agent](#).
- Determine the destination vaults for Hyper-V backups. See [Vaults for Hyper-V backups](#).
- Consider where to install Hyper-V Agent components to protect a Hyper-V standalone host or cluster. See [Recommended deployment for protecting a Hyper-V standalone host](#) and [Recommended deployment for protecting a Hyper-V cluster](#).

For best practices in a protected Hyper-V environment, see [Best practices in a protected Hyper-V environment](#).

For supported platform information, see the Hyper-V Agent release notes.

2.1 Portal for managing a Hyper-V Agent

The Hyper-V Agent is managed using Portal. You cannot manage the Hyper-V Agent using the legacy Windows CentralControl.

You must have a Portal account before you install the Hyper-V Agent. The account can be on a Portal instance that is hosted by your service provider, or installed on-premises.

If your Portal instance is installed on-premises, ensure that the Portal database is backed up so that the Hyper-V environment can be fully restored in the event of a disaster. Information for the Hyper-V Agent, including vault and backup job information, is saved in the Portal database. See [Recover jobs and settings from an offline Hyper-V Agent](#).

2.2 Vaults for Hyper-V backups

To provide fast, local vault access for backups and restores, back up Hyper-V data to an appliance or Satellite vault.

The data can then be replicated to a vault hosted by your service provider to ensure offsite protection in the case of a disaster.

If you choose not to use a Satellite vault, consider using a temporary vault to seed Hyper-V backups locally. The data can then be imported into an offsite vault.

2.3 Recommended deployment for protecting a Hyper-V standalone host

To protect a standalone Hyper-V host, we recommend the following:

- Install the Management service on a separate Windows server with local network access to the standalone host. The Management service server can be a virtual or physical machine that is on the

same domain as the Hyper-V standalone host.

- Install the Host service on the standalone host.

This deployment method minimizes the performance impact in the environment and avoids reboots on the standalone host when you install, upgrade or uninstall the Management service. However, if you do not want to install the Management service on a separate server, see [Alternate deployment for protecting a Hyper-V standalone host](#).

Note: You cannot install the Agent for Microsoft Windows on the standalone host. The Windows Agent is not compatible with the Host service.

2.3.1 Alternate deployment for protecting a Hyper-V standalone host

As described in [Recommended deployment for protecting a Hyper-V standalone host](#), we recommend installing the Management service on a separate virtual or physical Windows server with local network access to the standalone host, and installing the Host service on the standalone host.

However, if you do not want to install the Management service on a separate server, you can install both the Management service and Host service on a standalone host.

Beginning with Hyper-V Agent 8.84, this is not recommended. A driver required for file and folder restores is now installed with the Management service, and could require the host to be rebooted when you install, upgrade or uninstall the Agent Management service.

Note: You cannot install the Agent for Microsoft Windows on the standalone host. The Windows Agent is not compatible with the Host service.

2.4 Recommended deployment for protecting a Hyper-V cluster

To protect a Hyper-V cluster, we recommend the following:

- Install the Management service on a VM in the cluster, and enable High Availability on the VM. The VM where the Management service is installed must resolve to the same DNS server used by the Hyper-V cluster.
- Install the Host service on each host in the cluster. If the Host service is installed on all hosts in the cluster:
 - The Hyper-V Agent Management service automatically distributes the backup processing load across the hosts.

Beginning with Hyper-V Agent 8.84, after installing the Hyper-V Agent Management service, you must provide the Hyper-V environment network address and credentials in Portal before you can install Hyper-V Agent Host services. See [Configure a new Hyper-V Agent](#).

- Application-consistent backups can be created for VMs on all hosts. If the Host service is not installed on a host, backups for VMs on that host can only be crash-consistent.

If you do not want to deploy a VM in the cluster for the Management service, see [Alternate deployment for protecting a Hyper-V cluster](#).

Note: You cannot install the Host service on a host where the Agent for Microsoft Windows is installed. The Windows Agent is not compatible with the Host service.

2.4.1 Alternate deployment for protecting a Hyper-V cluster

As described in [Recommended deployment for protecting a Hyper-V cluster](#), we recommend installing the Hyper-V Agent Management service on a VM in the cluster, and installing the Host service on each host in the cluster.

If you do not want to deploy a VM in the cluster for the Management service, you can install the Management service on a supported Windows server that has local network access to the cluster. The server can be a physical or virtual machine that is on the same domain as the Hyper-V cluster.

Beginning with Hyper-V Agent 8.84, you cannot install the Management service directly on a Hyper-V host in a Hyper-V cluster. However, you can install the Management service on a standalone Hyper-V host.

You must install the Hyper-V Agent Host service on at least one host in a protected cluster. You do not have to install the Host service on every host in a cluster, since a single Host service can back up VMs on all hosts. However, this configuration is not optimal, for the following reasons:

- If the Host service is installed on only one host, all backup operations are delegated to the single host. The load cannot be distributed.
- A VM that is stored on a local volume can only be backed up if the Host service is installed on the host.
- Application-consistent backups cannot be created for VMs on a host where the Host service is not installed.
- A Hyper-V VM can only be restored to a host where the Host service is running. When restoring a Hyper-V VM in a cluster, you must choose a host where the Host service is running or the restore will fail.

If the Host service is not installed on the host where you want to restore a VM, you can restore the VM to a host where the Host service is installed, and then migrate the VM to the host that you want for the VM.

Note: You cannot install the Host service on a host where the Agent for Microsoft Windows is installed.

2.5 Hyper-V Rapid VM Restore requirements

Using Rapid VM Restore, you can restore a VM to a Hyper-V environment within minutes. You can then restore the VM permanently by migrating it to a permanent storage location in the Hyper-V environment. See [Restore a Hyper-V VM within minutes using Rapid VM Restore](#) and [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).

The following table lists and describes requirements for Hyper-V Rapid VM Restores. If the Agent, Portal and Vault requirements are not met, Rapid VM Restore does not appear as a restore option in Portal.

Component	Rapid VM Restore requirement
Hyper-V Agent	Hyper-V Agent version 9.00 or later. To perform a Rapid VM Restore in a Hyper-V cluster, the Hyper-V Agent Management server must be joined to the same domain as the cluster.
Portal	Portal version 8.89 or later.
Director Vault	A Director version 8.50 or later vault. The vault must be installed locally (i.e., not on a cloud server or in a remote datacenter). The Rapid VM Restore feature must be enabled on the vault. This feature is enabled by default on Satellite vaults. If you have a local Base vault, you can enable the Rapid VM Restore feature by running a script. See Enable the Rapid VM Restore feature on a vault .

2.5.1 Enable the Rapid VM Restore feature on a vault

To restore a VM within minutes using Rapid VM Restore, the VM backup must be saved in a local version 8.50 or later vault that has the Rapid VM Restore feature enabled.

The Rapid VM Restore feature is enabled by default on Satellite vaults. On Base vaults that are installed locally, you must enable the Rapid VM Restore feature using the following procedure.

To enable the Rapid VM Restore feature on a vault:

1. On the server where the vault is installed, open a PowerShell window as administrator, and navigate to the Scripts subfolder in the vault installation directory.
2. Run the following command:

```
.\VaultSettings.ps1 set IsRVMRAAllowed 1
```

2.6 Hyper-V Agent ports

The following table shows ports that must be open for the Hyper-V Agent to communicate with other systems:

Service	Port	Protocol	Communication
Management	Outbound: 8086, 8087	TCP	To Portal
	Outbound: 2546	TCP	To vault
	Outbound and inbound: 5444	TCP	With Host services
	Ports required for WMI connections. See documentation from Microsoft: Setting up a remote WMI connection	TCP	With cluster or standalone host
	Ports required for file sharing and WMI connections: 135-139 445	TCP/UDP TCP	With clients during file and folder restores

Service	Port	Protocol	Communication
Host	Outbound: 2546	TCP	To vault
	Outbound and inbound: 5444	TCP	With Management service

2.7 Best practices in a protected Hyper-V environment

For best performance, consider the following best practices for a Hyper-V environment that is protected by Hyper-V Agent 9.1:

- Enable the CSV Cache. In a failover cluster, enabling the CSV cache can improve Hyper-V Agent backup performance. Microsoft recommends enabling the CSV cache for read-intensive workloads. Search online for the following Microsoft documentation: *Use Cluster Shared Volumes in a Failover Cluster*

The CSV cache is enabled by default. To check that CSV cache is enabled, run the following PowerShell command:

```
Get-ClusterSharedVolume "csvName" | Get-ClusterParameter EnableBlockCache
```

- Avoid using VMs with limited support. On Windows Server 2016 or later, use VHDX format for virtual disks instead of VHD format. On Windows Server 2016 or later, Hyper-V Agent 9.1 does not back up VMs with disks in VHD format.

2.7.1 Best practices in Hyper-V on Windows Server 2012 R2

In Hyper-V on Windows Server 2016 or later, Hyper-V Agent 9.1 backs up VMs using features that are not available in Windows Server 2012 R2. In Hyper-V on Windows Server 2012 R2, Hyper-V Agent 9.1 uses the same backup method as previous agent versions.

The following best practices only apply in Hyper-V on Windows Server 2012 R2:

- Clean up checkpoints and snapshots before backups. On Windows Server 2012 R2, the Hyper-V Agent backs up and restores user-level checkpoints or snapshots with VMs, which can take a significant amount of time.

Note: On Windows Server 2016 or later, Hyper-V Agent 9.1 does not back up checkpoints.

Consistent with Microsoft best practices, we recommend not taking user-level snapshots or creating checkpoints of VMs that will be backed up in a production environment, except in a transient fashion. When it is necessary to take a snapshot or create a checkpoint of a protected VM, remove the snapshot or checkpoint before the next backup. Search online for the following Microsoft documentation: *Avoid using snapshots on a virtual machine that runs a server workload in a production environment*

- Use fixed-size virtual disks. If a VM includes a dynamically expanding virtual hard disk (VHDX or VHD), an incremental backup might be as large as a seed backup.

Note: On Windows Server 2016 or later, Hyper-V Agent 9.1 does not back up VMs with disks in VHD format.

- Avoid using VMs with limited or no backup support. On Windows Server 2012 R2, the Hyper-V Agent has limited support for VMs that contain:
 - Virtual disks which are configured as dynamic disks by Windows Disk Management (within a VM)
 - FAT or FAT32 volumes
 - Linux guest OS
 - No Hyper-V Integration Services running

During a backup, Hyper-V puts these VMs into a saved state for a brief period of time while capturing a VSS snapshot. The backup will be crash-consistent (not application-consistent).

During a backup, the Hyper-V Agent skips VMs that contain mixed storage or share virtual hard disks.

The Hyper-V Agent cannot back up a VM with 50 or more checkpoints. Microsoft specifies a maximum of 50 checkpoints for a VM. Search online for the following Microsoft documentation:

Maximums for virtual machines

3 Install, upgrade and uninstall the Hyper-V Agent

Before you can protect a Microsoft Hyper-V environment, you must:

- Install the Hyper-V Agent Management service. See [Install the Hyper-V Agent Management service](#).
- In Portal, specify Hyper-V environment information and credentials so that the Agent can authenticate with the environment that you want to protect. See [Configure a new Hyper-V Agent](#).
Beginning with version 8.84 of the Hyper-V Agent, you must provide Hyper-V environment information and credentials in Portal before you can install Hyper-V Agent Host services.
- Install the Hyper-V Agent Host service. See [Install the Hyper-V Agent Host service](#).

Note: You do not have to install agent software on individual VMs in the Hyper-V environment.

If a previous version of the Hyper-V Agent is installed, you can upgrade the Agent. See [Upgrade the Hyper-V Agent](#).

To uninstall a Hyper-V Agent, see [Uninstall the Hyper-V Agent Management service](#) and [Uninstall the Hyper-V Agent Host service](#).

3.1 Install the Hyper-V Agent Management service

Install the Hyper-V Agent Management service on a VM or server that has local network access to a protected Hyper-V environment. See [Prepare for a Hyper-V Agent deployment](#).

You cannot install the Management service directly on a Hyper-V host in a Hyper-V cluster. However, you can install the Management service on a standalone Hyper-V host. For recommended deployment methods, see [Recommended deployment for protecting a Hyper-V standalone host](#) and [Recommended deployment for protecting a Hyper-V cluster](#).

After installing the Hyper-V Agent Management service, you must provide the Hyper-V environment network address and credentials in Portal before you can install Hyper-V Agent Host services. See [Configure a new Hyper-V Agent](#).

By default, the Management service communicates with Host services using port 5444. However, you can specify a custom port during the Management service installation. Ensure that the correct inbound port is open.

To install the Management service silently, see [Install the Hyper-V Agent Management service in silent mode](#).

Note: All Hyper-V Agent services run under the LocalSystem account. The account for the Hyper-V Agent cannot be changed.

Note: Beginning in Hyper-V Agent 9.00, the startup type for Hyper-V Agent services is Automatic (Delayed Start). The delayed service start allows the Agent to clean up files from VMs running using Rapid VM Restore if an Agent host restarts.

To install the Hyper-V Agent Management service:

1. On the server or VM where you want to install the Management service, double-click the Hyper-V Agent Management service installation kit.
2. In the language list, click the language for the Agent, and then click **OK**.

The installation wizard starts.

3. On the Welcome page, click **Next**.
4. On the License Agreement page, read the license agreement. Click **I accept the terms in the license agreement**, and then click **Next**.
5. On the Destination Folder page, do one of the following:
 - To install the Management service in the default location, click **Next**.
 - To install the Management service in another location, click **Change**. In the Change Current Destination Folder dialog box, browse to the new installation folder, or enter it in the **Folder name** box. Click **OK**. On the Destination Folder page, click **Next**.
6. On the Register Hyper-V Agent Management with Portal page, specify the following information:
 - In the **Network Address** box, type the host name or IPV4 address of the Portal for managing the Hyper-V Agent.

Note: We recommend specifying the Portal host name when you register an agent to Portal. If the Portal IP address changes in the future, DNS can handle the change and you will not have to manually register the agent to Portal again.

- In the **Port** box, type the port number for communicating with the Portal.
- In the **Username** box, type the name of the Portal user for managing the Hyper-V Agent.

After the Hyper-V Agent is installed, the Agent appears on the Computers page of the Portal for this user and other Admin users in the user's site.

- In the **Password** box, type the password of the specified Portal user.
7. Click **Next**.
 8. On the Configure Communication Port page, specify the port used to communicate with Hyper-V Agent Host services, and then click **Next**.

By default, the Management service communicates with Host services using port 5444. Ensure that this inbound port, or the custom communication port specified, is open.

9. On the Ready to Install the Program page, click **Install**.
10. On the Installation Completed page, click **Finish**.

You must configure the Hyper-V environment network address and credentials before you can install Hyper-V Agent Host services. See [Configure a new Hyper-V Agent](#).

3.1.1 Install the Hyper-V Agent Management service in silent mode

Note: Before installing the Management service in silent mode, be sure that the port for communicating with Hyper-V Agent Host services is not in use.

To install the Management service in silent mode, run the following command in the directory where the installation kit is located:

```
installKitName /quiet /S [/L<localeID>] /V"/qn /L*v ["logFileName\"] UIREG_
NETADDRESS=webUIAddress [UIREG_PORT=webUIportNumber] UIREG_USERNAME=webUIUser
UIREG_PASSWORD=webUIUserPassword [COORDINATOR_PORT=portNumber]
[INSTALLDIR="\installPath\"]"
```

Where *installKitName* is the name of the Hyper-V Agent Management service installation kit: Hyper-V_Agent_Management-x-xx-xxxx.exe. x-xx-xxxx represents the Agent version number.

The following table lists and describes command parameters.

Parameter	Description
<i>/L<localeID></i>	Optional. Specifies the language for installation log messages. Available <i>localeID</i> values are: <ul style="list-style-type: none"> • 1033 – English (United States). This is the default value. • 1036 – French (Standard) • 1031 – German • 1034 – Spanish
<i>\"logFileName\"</i>	Optional. Specifies the path and name of the installation log file. If the <i>logFileName</i> includes spaces, enclose the value in double quotation marks. Example: <i>\"C:\Logs\My Log.txt\"</i> If you do not specify a <i>logFileName</i> , the installation log is saved in the Windows installer default location (usually the user's temp directory).
<i>UIREG_NETADDRESS=webUIAddress</i>	Specifies the host name or IP address of the Portal for managing the Hyper-V Agent. Example: <i>UIREG_NETADDRESS=192.0.2.233</i> Specifying the host name is recommended. This will allow DNS to handle IP address changes.
<i>UIREG_PORT=webUIportNumber</i>	Optional. Specifies the port number used to communicate with Portal. Example: <i>UIREG_PORT=8086</i> If you do not specify a <i>webUIportNumber</i> , port 8086 is used for communicating with Portal.

Parameter	Description
<code>UIREG_USERNAME=webUIUser</code>	Specifies the name of the Portal user associated with the Hyper-V Agent. Example: <code>UIREG_USERNAME=user@site.com</code>
<code>UIREG_PASSWORD=webUIUserPassword</code>	Specifies the password of the specified Portal user. Example: <code>UIREG_PASSWORD=password1234</code>
<code>COORDINATOR_PORT=portNumber</code>	Optional. Specifies the port used to communicate with Hyper-V Agent Host services. Example: <code>COORDINATOR_PORT=5444</code> If you do not specify a port, port 5444 is used for communicating with Hyper-V Agent Host services.
<code>INSTALLDIR=\"installFolder\"</code>	Optional. Specifies the installation folder for the Management service, if you do not want to install the Management service in the default location. The installation folder must be enclosed in double quotation marks if there are spaces in the path. Example: <code>INSTALLDIR=\"c:\Program Files\Management Service\"</code> If you do not specify an installation folder, the Management service is installed in the default location.

For example, to install the Management service in silent mode, you could run the following command:

```
Hyper-V_Agent_Management-x-xx-xxxx.exe /quiet /S /L1033 /V"/qn /L*v
\"C:\logs\1.log\" UIREG_NETADDRESS=192.0.2.233 UIREG_USERNAME=user@site.com
UIREG_PASSWORD=password1234 UIREG_PORT=8086 INSTALLDIR=\"C:\Program
Files\Management Service\""
```

3.2 Configure a new Hyper-V Agent

After the Hyper-V Agent Management service is installed and registered with Portal, you must configure the agent by specifying:

- The fully qualified domain name (FQDN) or IP address of the Hyper-V cluster or standalone host that you want to protect.
- Credentials for authenticating with the Hyper-V environment. For a Hyper-V cluster, the user must be an Active Directory domain user with administrative rights and full control over the cluster. For a standalone host, the user can be a local or domain user with administrative rights.

IMPORTANT: Beginning with version 8.84 of the Hyper-V Agent, you must provide Hyper-V environment information and credentials in Portal before you can install Hyper-V Agent Host services.

When configuring a new Hyper-V Agent, you can also add vault settings. Vault settings provides vault information and credentials so that the Agent can back up data to and restore data from the vault.

To change the Hyper-V environment information and credentials, add vault settings, or perform other configuration tasks after the initial configuration, see [Configure a Hyper-V Agent](#).

To configure a new Hyper-V Agent:

1. On the navigation bar in Portal, click **Computers**.

The Computers page shows registered computers.

2. Find the computer that has the Hyper-V Agent Management service installed, and expand its view by clicking its row.

Before you provide Hyper-V credentials for the Agent, the name of the computer where the Management service is installed appears on the Computers page.

3. If a Configuration mode section appears, select **Configure a new Hyper-V Agent**, and then click **Continue**.

The Configuration mode section appears if offline Hyper-V Agents are available in the site. This section includes a **Recover a previous Hyper-V Agent** option for recovering the configuration and backup jobs from an offline Hyper-V Agent instead of configuring a new Agent. See [Recover jobs and settings from an offline Hyper-V Agent](#).

4. In the Register agent with Hyper-V environment section, specify the following information:

- In the **Address** box, type the FQDN or IP address of the Hyper-V cluster or standalone host that you want to protect. Specifying the FQDN of the cluster or standalone host is recommended. This will allow DNS to handle IP address changes.

IMPORTANT: For a Hyper-V cluster, enter the FQDN or IP address of the cluster (not of a host in the cluster).

- In the **Domain** box, type the domain of the account for authenticating with the Hyper-V cluster or standalone host.

The domain is not required if you specify the domain in the **Username** box or if you specify a local user for a standalone host.

- In the **Username** box, type the administrator account that is used to authenticate with the Hyper-V cluster or standalone host. You can type the account as *username*, *domain\username*, or *username@domain*.

For a Hyper-V cluster, the user must be an Active Directory domain user with administrative rights and full control over the cluster.

For a standalone host, the user can be a local or domain user with administrative rights.

- In the **Password** box, type the password for the specified user.

5. To validate the credentials, click **Verify Information**. If the credentials are valid, a Success message appears. Click **Okay**.
6. Click **Continue**.

7. In the Vault Configuration section, click **Configure Vault**.

You can also add vault connections after the initial configuration. See [Add vault settings](#).

8. On the Vault Settings tab, click **Add Vault**.

9. In the Vault Settings dialog box, do one of the following:

- In the **Vault Name** box, enter a name for the vault connection. In the **Address** box, enter the vault host name or IPV4 address. In the **Account**, **Username**, and **Password** boxes, enter an account and credentials for backing up data to and restoring data from the vault.

Specifying the host name of the vault is recommended. This will allow DNS to handle IP address changes.

- If a policy with a vault profile is assigned to the computer, click the **Vault Profile** list. In the list, click the vault profile that you want to add for the computer. Vault information and credentials are then populated in the Vault Settings dialog box.

10. (Optional) Change one or more of the following Advanced Settings for the vault connection:

- **Agent Host Name**. Name of the computer on the vault. For a Hyper-V environment, by default, the name is the fully qualified domain name of the cluster or standalone host.
- **Port Number**. Port used to connect to the vault.
- **Attempt to Reconnect Every**. Specifies the number of seconds after which the agent should try to connect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 30 to 1800 seconds.
- **Abort Reconnect Retries After**. Enter the number of minutes after which the agent should stop trying to reconnect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 60 to 720 minutes. If the Agent cannot connect to the vault successfully in the specified time, the backup or restore fails.

11. Click **Save**.

If the agent is protecting a Hyper-V cluster, the FQDN of the cluster now appears on the Computers page in Portal instead of the Management service computer name.

3.3 Install the Hyper-V Agent Host service

The Hyper-V Agent Host service is installed on one or more hosts in a protected Hyper-V environment. See [Prepare for a Hyper-V Agent deployment](#).

Before you install a Host service, be sure that:

- The Hyper-V Agent Management service is installed on a server with local network access to the Hyper-V environment.
- The Hyper-V environment network address and credentials are specified in Portal for the Hyper-V Agent Management service. See [Configure a new Hyper-V Agent](#).

You must provide Hyper-V environment information and credentials in Portal before you can install Hyper-V Agent Host services. When you install a Hyper-V Agent Host service, the Management service checks whether the host is associated with the Hyper-V environment that is specified for the Management service in Portal. If the Hyper-V host is not associated with the Hyper-V environment specified in Portal, the installation will not proceed.

- There is local network connectivity to the Management service and the correct port is open. During the installation, the Host service must be able to establish connection with the Management service.

Do not install the Host service on the same machine as Agent for Microsoft Windows. The installer does not enforce this coexistence constraint.

To install the Host service silently, see [Install the Hyper-V Agent Host service in silent mode](#).

Note: All Hyper-V Agent services run under the LocalSystem account. The account for the Hyper-V Agent cannot be changed.

Note: Beginning in Hyper-V Agent 9.00, the startup type for Hyper-V Agent services is Automatic (Delayed Start). The delayed service start allows the Agent to clean up files from VMs running using Rapid VM Restore if an Agent host restarts.

To install the Hyper-V Agent Host service:

1. Log in to the Hyper-V host where you want to install the Host service.
2. Double-click the Hyper-V Agent Host service installation kit.
3. In the language list, click the language for the Agent, and then click **OK**.

The installation wizard starts.

4. On the Welcome page, click **Next**.
5. On the License Agreement page, read the license agreement. Click **I accept the terms in the license agreement**, and then click **Next**.
6. On the Destination Folder page, do one of the following:
 - To install the Host service in the default location, click **Next**.
 - To specify another installation location, click **Change**. In the Change Current Destination Folder dialog box, browse to the new installation location, or enter a folder in the **Folder name** box. Click **OK**. On the Destination Folder page, click **Next**.
7. On the Connection with Hyper-V Agent Management service page, in the **Network Address** box, enter the fully qualified domain name (FQDN) or IP address of the Hyper-V Agent Management service that will assign work to the Host service.

Note: Specifying the FQDN of the Management service is recommended. This will allow DNS to handle IP address changes.

If an error message states that the Hyper-V Agent Management service is unreachable, check that the Hyper-V environment network address and credentials have been specified for the

Management service in Portal, and that the host where you are installing the Host service is associated with the specified Hyper-V environment. See [Configure a new Hyper-V Agent](#).

8. In the **Port** box, enter the port number for communicating with the Hyper-V Agent Management service.

By default, the Management service communicates with Host services using port 5444. However, a custom port might have been specified during the Management service installation.

9. Click **Next**.
10. Click **Install**.
11. On the Wizard Completed page, click **Finish**.

3.3.1 Install the Hyper-V Agent Host service in silent mode

Before you install a Host service, be sure that:

- The Hyper-V Agent Management service is installed on a server with local network access to the Hyper-V environment.
- The Hyper-V environment network address and credentials are specified in Portal for the Hyper-V Agent Management service. See [Configure a new Hyper-V Agent](#).

You must provide Hyper-V environment information and credentials in Portal before you can install Hyper-V Agent Host services. When you install a Hyper-V Agent Host service, the Management service checks whether the host is associated with the Hyper-V environment that is specified for the Management service in Portal. If the Hyper-V host is not associated with the Hyper-V environment specified in Portal, the installation will not proceed.

- There is local network connectivity to the Management service and the correct port is open. During the installation, the Host service must be able to establish connection with the Management service.

To install the Host service in silent mode, run the following command in the directory where the installation kit is located:

```
installKitName /quiet /S [/L<localeID>] /V"/qn /L*v [\"logFileName\"]  
HOST=managementServiceAddress [PORT=portNumber] [INSTALLDIR=\"installPath\"]"
```

Where *installKitName* is the name of the Hyper-V Agent Host service installation kit: Hyper-V_Agent_Host-*x-xx-xxxx*.exe. *x-xx-xxxx* represents the Agent version number.

The following table lists and describes command parameters.

Parameter	Description
<code>/L<localeID></code>	Optional. Specifies the language for installation log messages. Available <i>localeID</i> values are: <ul style="list-style-type: none"> • 1033 – English (United States). This is the default value. • 1036 – French (Standard) • 1031 – German • 1034 – Spanish
<code>\"logFileName\"</code>	Optional. Specifies the path and name of the installation log file. If the <i>logFileName</i> includes spaces, enclose the value in double quotation marks. Example: <code>\"C:\Logs\My Log.txt\"</code> If you do not specify a <i>logFileName</i> , the installation log is saved in the Windows installer default location (usually the user's temp directory).
<code>HOST= managementServiceAddress</code>	Specifies the fully qualified domain name (FQDN) or IP address of the Hyper-V Agent Management service that assigns work to the Host service. Example: <code>HOST=192.0.2.234</code> Specifying the FQDN is recommended. This will allow DNS to handle IP address changes.
<code>PORT=<i>portNumber</i></code>	Optional. Specifies the port number for communicating with the Hyper-V Agent Management service. Example: <code>UIREG_PORT=5444</code> If you do not specify a port number, port 5444 is used.
<code>INSTALLDIR=<i>\"installFolder\"</i></code>	Optional. Specifies the installation folder for the Host service, if you do not want to install the Host service in the default location. The installation folder must be enclosed in double quotation marks if there are spaces in the name or path. Example: <code>INSTALLDIR=<i>\"c:\Program Files\Host Service\"</i></code> If you do not specify an installation folder, the Host service is installed in the default location.

For example, to install the Host service in silent mode, you could run the following command:

```
Hyper-V_Agent_Host-x-xx-xxxx.exe /quiet /S /L1036 /V"/qn /L*v \"C:\logs\1.log\"
HOST=192.0.2.234 PORT=5444"
```

3.4 Upgrade the Hyper-V Agent

To upgrade a Hyper-V Agent, first upgrade the Management service and then upgrade all Host services in the Hyper-V environment. See [Upgrade the Hyper-V Agent Management service](#) and [Upgrade the Hyper-V Agent Host service](#).

Note: All Hyper-V Agent services must be upgraded to the same version. Earlier service versions cannot be used with later service versions.

In a Hyper-V environment on Windows Server 2016 or later, incremental backups will reseed in some cases after an upgrade from Hyper-V Agent version 9.00 or earlier to version 9.1x:

- Backups will reseed for VMs with dynamically-expanding disks. Data will not be deduplicated on the vault after the reseed, and data from the previous Hyper-V Agent version will not be removed from the vault until specified by the retention settings. For example, if you use Hyper-V Agent 9.00 to back up a VM with dynamically-expanding disks and the resulting safeset is 25 GB in size, then upgrade the Hyper-V Agent to version 9.1x and back up the same VM (with no data changes) again, the next safeset will also be 25 GB in size and the pool size will increase to 50 GB.

IMPORTANT: After an upgrade to Hyper-V Agent 9.1x, the first backup of a VM with dynamically-expanding disks will be a full backup and may cause temporary billing overages or vault license exhaustion depending on your contract type. If you encounter this issue, please contact Support.

- Backups will partially reseed for VMs with fixed disks and user checkpoints.
- Backups will not reseed for VMs with fixed disks and no user checkpoints.

You can also move to a newer Agent version when recovering a protected Hyper-V environment after a disaster or when moving to a new Hyper-V environment. See [Recover jobs and settings from an offline Hyper-V Agent](#).

3.4.1 Upgrade the Hyper-V Agent Management service

Before upgrading the Management service, make sure that no backups or restores are running, and that the log viewer is not running.

IMPORTANT: You cannot upgrade the Management service to version 8.84 or later if it is installed directly on a Hyper-V host in a Hyper-V cluster. Instead, follow the procedure in [Replace a Hyper-V Agent Management service that is installed on a Hyper-V host](#). You can upgrade the Management service to version 8.84 or later if it is installed on a standalone Hyper-V host, but we recommend replacing it.

After upgrading the Management service, upgrade any Host services to the same version. See [Upgrade the Hyper-V Agent Host service](#).

To upgrade the Management service silently, see [Upgrade the Hyper-V Agent Management service in silent mode](#).

To upgrade the Hyper-V Agent Management service:

1. On the server or VM where you want to upgrade the Management service, double-click the Hyper-V Agent Management service installation kit.
2. In the confirmation dialog box, click **Yes**.
A message box warns you to be sure that there are no backups or restores in progress.
3. In the message box, click **Yes**.

4. In the installation wizard, click **Next**.
5. On the Installation Completed page, click **Finish**.

3.4.1.1 Upgrade the Hyper-V Agent Management service in silent mode

To upgrade the Management service in silent mode, run the following command in the directory where the installation kit is located:

```
installKitName /quiet /S [/L<localeID>] /V"/qn /L*v ["logFileName\"] "
```

Where *installKitName* is the name of the Hyper-V Agent Management service installation kit: Hyper-V_Agent_Management-x-xx-xxxx.exe. x-xx-xxxx represents the Agent version number.

The following table lists and describes command parameters.

Parameter	Description
\"logFileName\"	Optional. Specifies the path and name of the installation log file. If the logFileName includes spaces, enclose the value in double quotation marks. Example: \"C:\Logs\My Log.txt\" If you do not specify a logFileName, the installation log is saved in the Windows installer default location (usually the user's temp directory).

For example, to upgrade the Management service in silent mode, you could run the following command:

```
Hyper-V_Agent_Management-x-xx-xxxx.exe /quiet /S /L1033 /V"/qn /L*v  
\"C:\logs\1.1log\" "
```

3.4.1.2 Replace a Hyper-V Agent Management service that is installed on a Hyper-V host

You cannot upgrade the Management service to version 8.84 or later if it is installed directly on a Hyper-V host in a Hyper-V cluster. Instead, you must install a new Hyper-V Agent Management service on another VM or server and recover jobs and settings from the previous Management service version.

If the Management service is installed on a standalone Hyper-V host, you can upgrade the Management service to version 8.84 or later. However, we recommend replacing it with a Management service on another VM or server.

To replace a Hyper-V Agent Management service that is installed on a Hyper-V host:

1. Back up Hyper-V Agent logs in the <ManagementServiceInstallFolder>\Data folder.

This folder includes logs from both the Hyper-V Agent Management and Host services. Host services upload logs to the Management service after a process ends.
2. Uninstall the Management service that is installed on a Hyper-V host in a Hyper-V cluster.

3. Install the Hyper-V Agent Management service on a VM or server that has local network access to the protected Hyper-V environment. See [Prepare for a Hyper-V Agent deployment](#).
4. Recover jobs and settings from the offline Hyper-V Agent which you uninstalled in [Step 2](#). See [Recover jobs and settings from an offline Hyper-V Agent](#). Be sure to enter all required passwords, including Hyper-V environment, vault, and encryption passwords.

3.4.2 Upgrade the Hyper-V Agent Host service

Before upgrading the Host service, make sure that no backups or restores are running, that the log viewer is not running, and that the Management service has been upgraded to the same version. See [Upgrade the Hyper-V Agent Management service](#).

To upgrade the Host service silently, see [Upgrade the Hyper-V Agent Host service in silent mode](#).

To upgrade the Hyper-V Agent Host service:

1. On the server where you want to upgrade the Host service, double-click the Hyper-V Agent Host service installation kit.
2. In the confirmation dialog box, click **Yes**.
A message box warns you to be sure that there are no backups or restores in progress.
3. In the message box, click **Yes**.
4. In the installation wizard, click **Next**.
5. On the Installation Completed page, click **Finish**.

3.4.2.1 Upgrade the Hyper-V Agent Host service in silent mode

To upgrade the Host service in silent mode, run the following command in the directory where the installation kit is located:

```
installKitName /quiet /S [/L<localeID>] /V"/qn /L*v ["logFileName"]  
HOST=managementServiceAddress "
```

Where *installKitName* is the name of the Hyper-V Agent Host service installation kit: Hyper-V_Agent_Host-x-xx-xxxx.exe. x-xx-xxxx represents the Agent version number.

The following table lists and describes command parameters.

Parameter	Description
<code>"/logFileName"</code>	Optional. Specifies the path and name of the installation log file. If the <i>logFileName</i> includes spaces, enclose the value in double quotation marks. Example: <code>"/C:\Logs\My Log.txt"</code> If you do not specify a <i>logFileName</i> , the installation log is saved in the Windows installer default location (usually the user's temp directory).

For example, to upgrade the Host service in silent mode, you could run the following command:

```
Hyper-V_Agent_Host-x-xx-xxxx.exe /quiet /S /L1036 /V"/qn /L*v \"C:\logs\1.log\""
```

3.5 Uninstall the Hyper-V Agent Management service

Note: If you reinstall the Management service in a Hyper-V environment for any reason, you must also reinstall each Host service.

To uninstall the Hyper-V Agent Management service:

1. On the machine where you want to uninstall the Management service, double-click the Hyper-V Agent Management service installation kit.
2. On the Welcome page, click **Next**.
3. A warning message asks you to ensure that there are no backups or restores in progress. To proceed with uninstalling the Management service, click **Yes**.
4. On the Remove the Program page, click **Remove**.
5. If the Files in Use page appears, select **Automatically close and attempt to restart applications**, and then click **OK**.
6. On the InstallShield Wizard Completed page, click **Finish**.

3.5.1 Uninstall the Management service in silent mode

To uninstall the Management service in silent mode, run the following command:

```
installKitName /quiet /S [/L<localeID>] /V"/qn REMOVE=ALL /L*v [\"logFileName\"]
```

Where *installKitName* is the name of the Hyper-V Agent Management service installation kit: Hyper-V_Agent_Management-x-xx-xxxx.exe. x-xx-xxxx represents the Agent version number.

The following table lists and describes command parameters.

Parameter	Description
[/L<localeID>]	Optional. Specifies the language for uninstallation log messages. Available <i>localeID</i> values are: <ul style="list-style-type: none"> • 1033 – English (United States). This is the default value. • 1036 – French (Standard) • 1031 – German • 1034 – Spanish

Parameter	Description
<code>"logFileName"</code>	<p>Optional. Specifies the path and name of the uninstallation log file. If the logFileName includes spaces, enclose the value in double quotation marks.</p> <p>Example: <code>"C:\Logs\My Log.txt"</code></p> <p>If you do not specify a logFileName, the uninstallation log is saved in the Windows installer default location (usually the user's temp directory).</p>

3.6 Uninstall the Hyper-V Agent Host service

Do not uninstall the Host service when a backup or restore operation is running. Uninstalling the Agent Host service during any operation will leave the Agent in an inconsistent state.

You can later reinstall an Agent Host service without having to restore state information, as long as no backup or restore operations are in progress.

To uninstall the Hyper-V Agent Host service:

1. On the Hyper-V host where you want to uninstall the Host service, double-click the Hyper-V Agent Host service installation kit.
2. On the Welcome page, click **Next**.
3. A warning message asks you to ensure that there are no backups or restores in progress. To proceed with uninstalling the Host service, click **Yes**.
4. On the Remove the Program page, click **Remove**.
5. If the Files in Use page appears, select **Automatically close and attempt to restart applications**, and then click **OK**.
6. On the InstallShield Wizard Completed page, click **Finish**.

3.6.1 Uninstall the Host service in silent mode

To uninstall the Host service in silent mode, run the following command in the directory where the installation kit is located:

```
installKitName /quiet /S /V"/qn REMOVE=ALL"
```

Where *installKitName* is the name of the Hyper-V Agent Host service installation kit: Hyper-V_Agent_Host-x-xx-xxxx.exe. x-xx-xxxx represents the Agent version number.

4 Configure a Hyper-V Agent

You must provide Hyper-V environment information and credentials when installing the agent. You can also add vault settings. See [Configure a new Hyper-V Agent](#).

After the agent is installed, you can:

- [Change credentials or the network address for accessing Hyper-V](#).
- [Add vault settings](#). Vault settings provide vault information and credentials so that the agent can back up data to and restore data from the vault.
- Add a description for the agent. The description appears for the Hyper-V environment on the Computers page. See [Add a description](#).
- Add retention types that specify how long backups are kept on the vault. See [Add retention types](#).
- Configure email notifications so that users receive emails when backups complete, fail, or have errors. See [Set up email notifications for backups on a computer](#).
- Specify the amount of bandwidth consumed by backups. See [Configure bandwidth throttling](#).
- Resolve a certificate failure reported by a Hyper-V Agent. See [Resolve certificate failures](#).

4.1 Change credentials or the network address for accessing Hyper-V

To change credentials or the network address for accessing Hyper-V:

1. On the navigation bar in Portal, click **Computers**. The Computers page shows registered computers. Find the Hyper-V Agent for which you want to change Hyper-V credentials, and click the computer row to expand its view.
2. Click the **Advanced** tab.
3. On the **Cluster Credentials** tab, specify the following information:
 - In the **Address** box, type the host name or IP address of the Hyper-V cluster or standalone host that you want to protect. Specifying the host name of the cluster or standalone host is recommended. This will allow DNS to handle IP address changes.

IMPORTANT: For a Hyper-V cluster, enter the host name or IP address of the cluster (not of a host in the cluster).
 - In the **Domain** box, type the domain of the account for authenticating with the Hyper-V cluster or standalone host.

The domain is not required if you specify the domain in the **Username** box or if you specify a local user for a standalone host.
 - In the **Username** box, type the administrator account that is used to authenticate with the Hyper-V cluster or standalone host. You can type the account as *username*, *domain\username*, or *username@domain*.

For a Hyper-V cluster, the user must be an Active Directory domain user with administrative rights and full control over the cluster.

For a standalone host, the user can be a local or domain user with administrative rights.

- In the **Password** box, type the password for the specified user.
4. To validate the credentials, click **Verify Information**. If the credentials are valid, a message appears. Click **Okay**.
 5. Click **Save**.

4.2 Add vault settings

Before a computer can back up data to or restore data from a vault, vault settings must be added for the computer. Vault settings provide vault information, credentials, and agent connection information required for accessing a vault.

When adding vault settings for a computer, Admin users and regular users can manually enter vault information, or select a vault profile with vault information and credentials.

If a policy is assigned to a computer, Admin users can select any vault profile from the policy. Regular users can only select policy vault profiles that are also assigned to them.

If a policy is not assigned to a computer, Admin users can select any vault profile in the site. Regular users can only select vault profiles that are assigned to them.

You can also add a vault connection during the initial Hyper-V Agent configuration. See [Configure a new Hyper-V Agent](#).

Over-the-wire encryption is automatically enabled when you add vault settings or save existing vault settings.

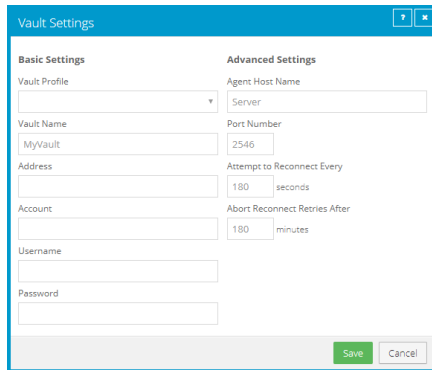
To add vault settings:

1. On the navigation bar in Portal, click **Computers**.
2. Find the computer for which you want to add vault settings, and click the computer row to expand its view.

If the Configure Manually box appears, click **Configure Manually**. The Configure Manually box appears for some computers where a backup job has not been created.

3. On the Vault Settings tab, click **Add Vault**.

The Vault Settings dialog box appears.



4. Do one of the following:

- In the **Vault Name** box, enter a name for the vault. In the **Address** box, enter the vault host name or IPV4 address. In the **Account**, **Username**, and **Password** boxes, enter an account and credentials for backing up data to and restoring data from the vault.

Specifying the host name of the vault is recommended. This will allow DNS to handle IP address changes.

- Click the **Vault Profile** box. If one or more vault profiles appear, click the vault profile that you want to add for the computer. Vault information and credentials are then populated in the Vault Settings dialog box.

If a policy is assigned, the **Vault Profile** list includes vault profiles from the policy. If a policy is not assigned, the list includes vault profiles from the site. For a regular user, the list only includes vault profiles that are also assigned to the user.

5. (Optional) Change one or more of the following Advanced Settings for the vault connection:

- **Agent Host Name.** Name to use for the computer on the vault. For a Hyper-V environment, by default, the name is the fully qualified domain name of the cluster or standalone host.
- **Port Number.** Port used to connect to the vault. The default port is 2546.
- **Attempt to Reconnect Every.** Specifies the number of seconds after which the agent should try to connect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 30 to 1800 seconds.
- **Abort Reconnect Retries After.** Enter the number of minutes after which the agent should stop trying to reconnect to the vault if the vault becomes unavailable during a backup or restore. The value can be from 60 to 720 minutes. If the Agent cannot connect to the vault successfully in the specified time, the backup or restore fails.


6. Click **Save**.

4.3 Add a description

You can add a description for a computer in Portal. The description appears on the Computers page, and can help you find and identify a particular computer.

To add a description:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to add a description, and click the row to expand its view.
If the Configure Manually box appears, click **Configure Manually**. The Configure Manually box appears for some computers where a backup job has not been created.
3. On the Advanced tab, click the **Options** tab.
4. In the Agent Description box, enter a description for the computer.



5. Click **Save**.

4.4 Add retention types

When you schedule or run a backup job, you must select a retention type for the resulting safeset. A retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.

Portal Admin users and regular users can add retention types for a computer where a policy is not assigned.

If a policy is assigned to a computer, retention types cannot be added or modified on the Computers page. Instead, retention types can only be added or modified in the policy.

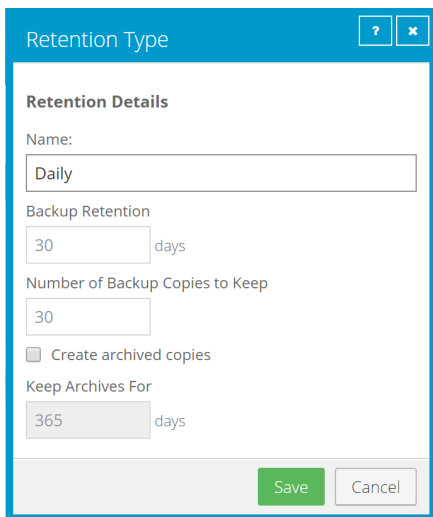
To add a retention type:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to add a retention type, and click the row to expand its view.
If the Configure Manually box appears, click **Configure Manually**. The Configure Manually box appears for some computers where a backup job has not been created.
3. On the Advanced tab, click the **Retention Types** tab.

If a policy is assigned to the computer, you cannot add or change values on the Retention Types tab. Instead, retention types can only be added or modified in the policy.

4. Click **Create Retention Type**.

The Retention Type dialog box appears.



5. Complete the following fields:

Name	Specifies a name for the retention type.
Backup Retention	Specifies the number of days a safeset is kept on the vault. A safeset is deleted when its expiry date is reached. <i>Note:</i> Safesets are not deleted unless the specified number of copies online has also been exceeded.
Number of Backup Copies to Keep	Specifies how many safesets from a backup job are stored online. It functions in a first in/first out manner. Once the number of safesets is exceeded, the oldest safesets are automatically deleted until the actual number of safesets matches the definition. <i>Note:</i> Safesets are not deleted unless the specified number of days online has also been exceeded.
Create archived copies	Select this check box to create archived copies of safesets.

Keep Archives For	<p><i>Note:</i> If data archiving is disabled in your Portal instance, this value does not appear.</p> <p>Specifies how long the data is stored offline. Archive storage is used to store data offline for long periods of time. This data is not immediately accessible since it is stored remotely. A greater amount of time is required to restore from archive media. Typically, only long-term data archives are stored offline. The parameters for archived data are from 365 to 9999 days.</p> <p>Assuming that at least one successful backup has been completed for the job, there will never be less than one copy of its backup online. This is true even if all retention settings are zero, expiry conditions have been met, and the job definition is deleted from your system. Deleting the job does not affect data on the vault. Only your service provider can remove jobs and their associated data from the vault. This is a safeguard against accidental or malicious destruction of data.</p>
-------------------	---

6. Click **Save**.

4.5 Set up email notifications for backups on a computer

Email notifications for backups in Hyper-V environments are configured separately for each Hyper-V Agent.

To set up email notifications for a computer:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the computer for which you want to configure email notifications, and click the computer row to expand its view.
3. On the **Advanced** tab, click the **Notifications** tab.

If the Notifications tab appears, but a policy is assigned to the computer, you cannot change values on the Notifications tab. Instead, notifications can only be modified in the policy.

4. Select one or more of the following checkboxes:
 - **On failure.** If selected, users receive an email notification when a backup or restore fails. If a backup fails, you cannot recover any files from the backup.

- **On error.** If selected, users receive an email notification when a backup or restore completes with errors in the log file. You cannot recover files that are backed up with errors, but you can restore other files from the backup (safeset).
- **On successful completion.** If selected, users receive an email notification when a backup or restore completes successfully. You can recover files from a backup that completes, even if there are warnings in the log file.

Email notifications are sent separately for each backup and restore. For example, if three backup jobs fail on a computer and **On failure** is selected for the computer, three notification emails are sent.

If users will receive email notifications after backups and restores, specify the following email notification information:

Email "From" Address	Email address from which email notifications will be sent.
Outgoing Mail Server (SMTP)	Network address of the SMTP that will send the email.
Recipient Address(es)	Email notification recipient email addresses, separated by commas. These should be real, valid email addresses. If one or more is not valid, the transmission to those addresses will fail, and errors will appear in the log files.
Outgoing Server Port (SMTP)	Port number for sending email notifications.
SMTP Credentials	If required, SMTP username, domain, and password.

5. Click **Save**.

4.6 Configure bandwidth throttling

For the Hyper-V Agent, bandwidth throttling is applied at the Host level, and applies only to backups. The total bandwidth sent to the vault can be as high as the specified maximum multiplied by the number of nodes where the Host service is installed.

Bandwidth settings include:

- Maximum bandwidth (upper limit), in megabits per second, to be consumed by the Agent for backups and restores.
- Period of time during the day that throttling is in effect. Only one time window can be specified. Outside the window, no throttling takes place.
- Days of the week that throttling is in effect.

If the bandwidth throttling time period begins when a backup is underway, the maximum bandwidth is applied dynamically to the running backup. Similarly, if the bandwidth throttling time period ends when a backup is running, bandwidth throttling is ended for the backup.

If you edit a computer’s bandwidth settings while a backup is running, the new settings do not affect the backup that is running. Bandwidth settings are applied when a backup starts, and are not applied to backups that are already running.

If a policy is assigned to a computer, bandwidth throttling settings cannot be modified on the Computers page. Instead, settings can only be added or modified in the policy.

To configure bandwidth throttling:

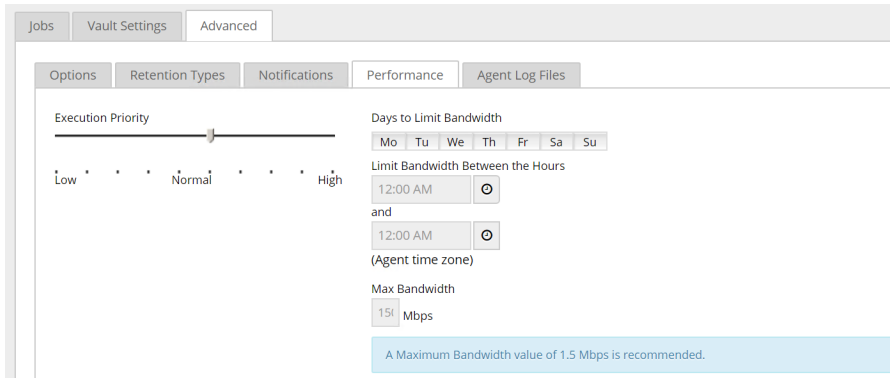
1. On the navigation bar, click **Computers**.
2. Find the computer for which you want to configure bandwidth throttling, and click the row to expand its view.

If the Configure Manually box appears, click **Configure Manually**. The Configure Manually box appears for some computers where a backup job has not been created.

3. Click the **Advanced** tab, click the **Performance** tab, and then edit the bandwidth settings.

If a policy is assigned to the computer, you cannot add or change values on the Performance tab. Instead, bandwidth settings can only be modified in the policy.

Note: Depending on your Internet speed, the recommended maximum bandwidth value (1.5 Mbps) shown in Portal may be low. This is only a recommendation. You can specify a higher maximum bandwidth if your Internet speed will support it.



4. Click **Save**.

4.7 Resolve certificate failures

If an agent reports a certificate failure, you must resolve the failure before backups and restores can continue. Certificate failures are summarized in the Current Snapshot on the Dashboard and shown on the Computers page in Portal. See [Monitor backups and computers using the Current Snapshot](#) and [View computer and job status information](#). Agents can report certificate failures if they support certificate pinning, a security feature that is designed to ensure that agents are connecting to legitimate vaults.

A certificate failure can occur when a Hyper-V agent tries to connect to a Director version 8.60 or later vault where certificate pinning is enabled. Beginning with Hyper-V Agent 9.00, when a Hyper-V agent tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault.

If a certificate failure is reported, please contact your IT security staff or service provider to determine whether the certificate change was expected or whether further investigation is required.

If the certificate change was expected, follow the steps below to re-pin the certificate. When you re-pin a certificate, the agent securely records the new public key of the certificate.

To resolve certificate failures:

1. On the navigation bar, click **Computers**. The Computers page shows registered computers.
2. Select the check box for each computer with a certificate failure that you want to resolve.

Note: Only select computers that have the Certificate failure status, or the Re-pin certificate action will not be available.

3. In the **Actions** list, click **Re-pin certificate**.
4. In the confirmation dialog box, click **Yes**.
5. In the Success message box, click **Okay**.

5 Add and schedule a Hyper-V backup job

After a Hyper-V environment is added in Portal, you can create a backup job that protects VMs in the cluster or standalone host. The backup job specifies virtual machines (VMs) to back up, specifies where to save the backup data, and includes schedules for running the backup job.

Beginning with version 8.84 of the Hyper-V Agent, you can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows VMs in Hyper-V environments. Application-consistent backups minimize the amount of work needed to restore applications from backups. You can also specify whether application transaction logs should be truncated during application-consistent backups.

Each VM in a Hyper-V environment can only be included in one backup job at a time. If a VM is already included in a backup job, you cannot add it to another job.

For best practices when creating and running backup jobs, see [Best practices for backing up Hyper-V VMs](#). For best practices when seeding VM backups, see [Best practices for seeding Hyper-V VM backups](#). To create Hyper-V backup jobs, see [Add a Hyper-V backup job](#).

When you run a Hyper-V Agent backup job, each VM in the job is backed up as a separate job (task) on the vault, and is automatically assigned a unique task name. This differs from jobs created using traditional agents, where each backup job is associated with a single task on the vault. This Hyper-V Agent backup job design provides a number of benefits:

- VMs in a single job can be backed up concurrently.
- Backup processing for individual VMs can be distributed across multiple nodes in a Hyper-V cluster.
- The Hyper-V Agent is scalable in large Hyper V environments.
- A VM can be moved to another job without reseeding (if encryption credentials are the same in both jobs).
- A protected VM can be restored even if its backup job has been deleted, as long as the VM's backup has not been deleted from the vault.
- If a protected VM has been deleted from your environment, and is no longer associated with a backup job, you can still see the VM in Portal in the protected view, and restore the VM from the vault. After restoring the VM, you can add the VM to a new job with the same encryption password, and continue to back up the VM without reseeding.

To view the vault task name for each protected Hyper-V VM, see [Determine the name of a VM's task on the vault](#).

All Hyper-V backup data is protected using AES 256 encryption.

5.1 Best practices for backing up Hyper-V VMs

Consider the following best practices when creating and running backup jobs using Hyper-V Agent 9.1:

- Include more than one VM in a backup job
- Avoid creating a separate backup job for each VM. The agent is optimized for backing up multiple VMs concurrently in one job.
- Avoid unnecessary reseeds. After the first backup for a Hyper-V VM, the agent only sends data that has changed since the last backup to the vault. However, under some circumstances, “reseeds” can occur. In a reseed, all data for a VM is sent to the vault even though the VM was previously backed up.

The following list describes situations when backups reseed:

- VM backups in a job reseed if the job’s encryption password changes.
- A VM backup reseeds if it is backed up as part of one backup job, and then moved to a job with a different encryption password.

If a VM is moved to a different backup job with the same encryption password, the VM backup does not reseed.
- On Windows Server 2012 R2, snapshot (AVHD) files reseed after storage migration. Hard disk VHD(x) files do not reseed after storage migration.

5.1.1 Best practices in Hyper-V on Windows Server 2012 R2

In Hyper-V on Windows Server 2016 or later, Hyper-V Agent 9.1 backs up VMs using features that are not available in Windows Server 2012 R2. In Hyper-V on Windows Server 2012 R2, Hyper-V Agent 9.1 uses the same backup method as previous agent versions.

The following best practices only apply in Hyper-V on Windows Server 2012 R2:

- Where possible, include VMs on the same CSV in the same backup job.
- If multiple jobs are needed to back up VMs on the same set of CSVs, stagger the job schedules so that the jobs do not run at the same time.

5.2 Best practices for seeding Hyper-V VM backups

The first backup for a Hyper-V VM is a “seed” backup, in which all VM data is sent to the vault. Consider the following best practices when seeding Hyper-V VM backups.

- Seed backups locally
- Ideally, use a Satellite vault to provide fast, local vault access. If you do not use a Satellite vault, consider using a temporary vault to seed Hyper-V backups locally. The data can then be imported into an offsite vault.
- Seed VM backups in separate jobs

- Seeding backups, particularly for large VMs, can take a significant amount of time. Because deferring is not available for scheduled Hyper-V backup jobs, it is best not to seed VMs in a scheduled job. If you add a VM to an existing scheduled job, the job could take a long time, and potentially cause the backup to overlap the next scheduled backup.

To seed a VM backup, we recommend creating a temporary job with the VM. You can seed the VM backup by running the temporary job manually (ad hoc) with deferring, and then move the VM to an existing scheduled job. To avoid reseeding, the encryption password and vault must be the same in the temporary job and in the job where you eventually add the VM.

Note: Normally, it is best to include multiple VMs in a single backup job. See [Best practices for backing up Hyper-V VMs](#).

To create and run a job manually (ad hoc) to seed a VM backup:

1. Create a temporary backup job that includes the VM that you want to seed. Ensure that the encryption password and vault for the temporary job are the same as the password and vault for the job where you eventually want to add the VM. See [Add a Hyper-V backup job](#).
2. Run the temporary job manually (ad hoc). You can enable deferring when running the job manually, and run the job multiple times until the VM backup is completely seeded. See [Run an ad-hoc backup](#).
3. After the VM backup is seeded, move the VM out of the temporary job, and add it into an existing scheduled job. See [Edit a Hyper-V backup job](#).

The VM backup will continue without reseeding because the vault and encryption password are the same in the temporary job and in the existing scheduled job.

5.3 Application-consistent backups on Hyper-V VMs

Beginning in version 8.84, the Hyper-V Agent can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows virtual machines (VMs) in Hyper-V environments.

Note: A Hyper-V Agent backup is not sufficient for an authoritative restore of Active Directory objects. For an authoritative restore, a System State backup with the Windows Agent is required.

In an application-consistent backup, pending application transactions are written to disk before the data is backed up. This minimizes the amount of work required to restore the application. If application-consistency is not enabled in a backup job, the backups are crash-consistent. In a crash-consistent backup, pending application transactions are rolled back and manual steps are required to ensure that applications are completely restored.

To create application-consistent backups on Hyper-V VMs, you must provide credentials that have admin access to VMs. You can provide guest VM credentials with admin access to multiple VMs in a backup job and/or provide credentials for specific VMs. If you provide credentials for a specific VM, the guest VM credentials for multiple VMs will never be used to connect to that VM.

If you enable application-consistent backups in a backup job but an application-consistent backup cannot be created for a VM, the Hyper-V Agent creates a crash-consistent backup for the VM. To check whether each VM backup is application-consistent or crash-consistent, view the backup log.

An application-consistent backup cannot be created in the following cases:

- The guest VM credentials are incorrect.
- The VM is not in a running state (e.g., is powered off, paused or migrating).
- The VM is running on a Hyper-V host with an operating system where application-consistent backups are not supported.
- The VM is running on a Hyper-V host where the Host service is not installed.
- The VM has a guest operating system where application-consistent backups are not supported.
- The VM configuration version is earlier than 6.2.
- The Backup (volume shadow copy) integration service is not enabled for the VM.
- A VSS writer on the host cannot connect to the VM.
- A VSS writer on the VM is not available or is in a bad state.

Note: To determine whether VSS writers required for an application-consistent backup are available on a VM (e.g., the SQL writer for a SQL Server backup), check the backup log. The backup log lists VSS writers on each VM and indicates the status of each writer.

For a list of supported Hyper-V host and guest operating systems, see the Hyper-V Agent release notes.

Log truncation in application-consistent backups

When performing application-consistent backups, the Hyper-V Agent can truncate SQL Server, Exchange and SQL transaction logs for SharePoint Server. This prevents the transaction logs from taking up a significant amount of disk space and reducing system performance. There are no logs to truncate when performing application-consistent backups of Domain Controllers with Active Directory.

Note: The Hyper-V Agent can truncate transaction logs for the default SQL Server instance and for all Exchange Server databases. The Hyper-V Agent cannot truncate logs for named SQL Server instances.

To truncate transaction logs on a VM after an application-consistent backup, you must enable log truncation in the backup job.

To check whether log truncation was successful on each VM after a backup, view the backup logs.

Note: If you also back up databases using another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

5.4 Add a Hyper-V backup job

After a Hyper-V environment is added in Portal, you can create a backup job that specifies which virtual machines (VMs) to back up, and where to save the backup data.

You must add vault settings and Hyper-V environment information before you can add a backup job. See [Configure a new Hyper-V Agent](#).


You can start to create a Hyper-V backup job by selecting VMs to include in the job. If you select a VM that is already included in another job, the VM will not be added to the new job. You can also start to create a Hyper-V backup job without selecting VMs on the Virtual Machines tab.

Beginning with version 8.84 of the Hyper-V Agent, you can create application-consistent backups of Microsoft SQL Server, Exchange, SharePoint and Active Directory on Windows VMs in Hyper-V environments. Application-consistent backups minimize the amount of work needed to restore applications from backups. You can also specify whether application transaction logs should be truncated during application-consistent backups. For more information, see [Application-consistent backups on Hyper-V VMs](#).

To add a Hyper-V backup job:

1. On the navigation bar, click **Computers**.

The Computers page shows registered computers and environments.

2. Click the Hyper-V environment row. 

3. Do one of the following:

- To start creating the job without selecting VMs to include in the job:

- a. Click the **Jobs** tab.
- b. In the **Select Job Task** menu, click **Create New Hyper-V Job**.

The Create New Job dialog box appears. The **Unprotected** box shows VMs that are not included in a backup job.

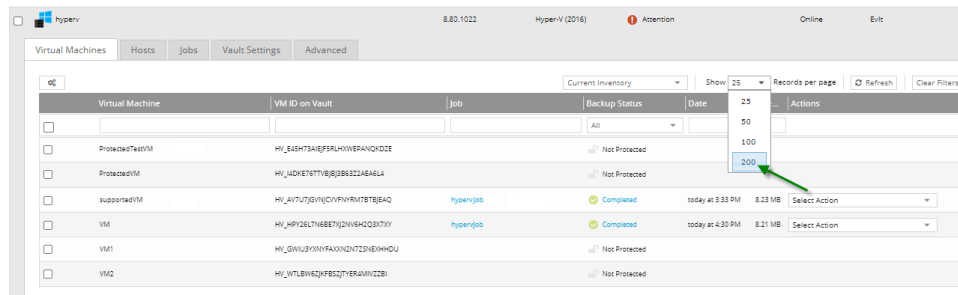
- To start creating the job by selecting VMs to include in the job:

- a. Click the **Virtual Machines** tab.

The Virtual Machines tab lists VMs in the Hyper-V environment. The Jobs column is blank for VMs that are not included in a backup job. By default, the Virtual Machines tab shows 25 VMs at a time.

- b. If the VMs that you want to include in the job are not listed on the tab, click the **Show <number of> records per page** list, and click the number of VMs to show.

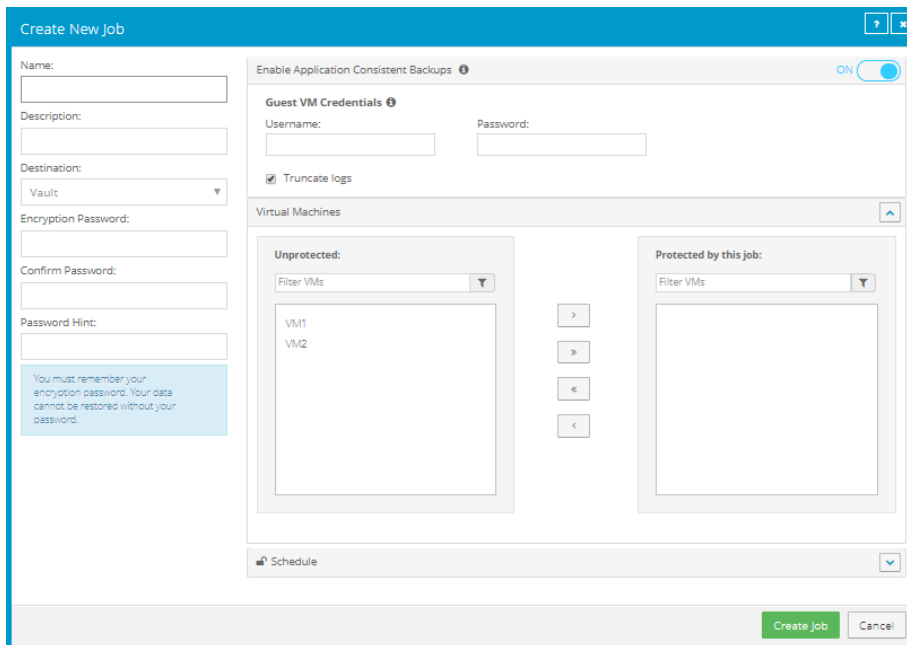
A maximum of 200 VMs can appear on the Virtual Machines tab at the same time, and you can only select VMs that appear on the Virtual Machines tab. However, you can add more VMs to the job later in this procedure.



- c. Do one of the following:
- Select the check box for each VM that you want to include in the backup job.
 - Select the check box at the top left of the list to select all VMs on the page.
- If you select a VM that is included in another job, it will not be added to the new job.

d. Click **Create Hyper-V Job**. 

The Create New Job dialog box appears. The **Unprotected** box shows VMs that are not included in a backup job. If you selected VMs on the Virtual Machines tab, the **Protected by this job** box shows the selected VMs.



4. In the Create New Job dialog box, specify the following information:
- In the **Name** box, type a name for the backup job.
 - In the **Description** box, type a description for the backup job.
 - In the **Destination** list, select the vault where you want to save the backup data.

- In the **Encryption Password** and **Confirm Password** boxes, enter a data encryption password. You can also enter a password hint in the **Password Hint** box.

IMPORTANT: You must enter the encryption password to recover your data. If you forget the password, you lose access to your data. The password is not maintained anywhere else and cannot be recovered.

Note: Hyper-V backup data is encrypted using the AES 256 encryption method.

5. In the Enable Application Consistent Backups box, do one of the following:

- To perform crash-consistent backups of VMs in the backup job, turn off the **Enable Application Consistent Backups** toggle.
- To perform application-consistent backups of SQL Server, Exchange, SharePoint, and Active Directory on any VMs in the backup job, do the following:
 - a. Turn on the **Enable Application Consistent Backups** toggle.
 - b. To enter credentials for VMs in the job, enter an admin user's username and password in the Guest VM Credentials area.

The specified user must have admin access to VMs in the backup job. You can enter the username as *username* or *domain\username*.

You can also enter credentials for specific VMs in the job. See [Step 7](#). If you enter credentials for a specific VM in the job, the Agent will not attempt to connect to the VM using the Guest VM Credentials.


IMPORTANT: If you do not enter credentials for VMs in a job where application-consistent backups are enabled, backups will be crash-consistent. Credentials are required for all application-consistent backups in Hyper-V environments— with or without log truncation.

- c. Do one of the following:
 - To preserve application transaction logs on VMs in the job, clear **Truncate logs**.
 - To truncate application logs on VMs in the job, select **Truncate logs**.


Note: If you also back up databases with another tool (e.g., native SQL Server backup), use only one tool for truncating logs.

6. Do one or more of the following until the **Protected by this job** box shows all VMs that you want to include in the job:

- To find one or more VMs in the **Unprotected** or **Protected by this job** box, enter characters from the VM names in the associated **Filter VMs** box.

- To add all VMs in the **Unprotected** box to the backup job, click **Protect all.** 

- To add some VMs in the **Unprotected** box to the backup job, select the VMs in the

Unprotected box, and then click **Protect selected.** 

To select multiple VMs in the list, press CTRL and click the VM names. To select multiple consecutive VMs in the list, press Shift and then click the first and last VM that you want to select, or drag the mouse across the VMs.

- To remove all VMs in the **Protected by this job** box from the backup job, click **Unprotect all**.



- To remove some VMs in the **Protected by this job** box from the backup job, select the VMs in the **Protected by this job** box, and then click **Unprotect selected**.




To select multiple VMs in the list, click each VM name.

Ensure that each VM that you want to include in the backup job appears in the **Protected by this job** box.

7. If application-consistent backups are enabled in the backup job and you want to enter credentials for a specific VM in the job, click the arrow at the right side of the VM name in the **Protected by this job** box, and enter an admin user's username and password for the VM.

The specified user must have admin access to the VM. You can enter the username as *username* or *domain\username*.



Note: If you enter credentials for a specific VM in the job, the agent will not attempt to connect to the VM using the Guest VM Credentials.

8. To schedule the backup job to run, click **Schedule**. In the **Schedule** box that appears, create one or more schedules. See [Add or edit a schedule for a Hyper-V backup job](#).
9. Click **Create Job**.

5.5 Edit a Hyper-V backup job

You can edit an existing Hyper-V backup job to change one or more of the following:

- VMs that are included in the job
- Encryption password and password hint
- Passwords for application-consistent backups (e.g., if VM passwords change in accordance with your organization's password change requirements)
- Schedules and retention types


You cannot change a backup job's name or vault connection.

Because each VM is backed up as a separate safeset on the vault, you can move a VM from one backup job to another without causing the backup to reseed. As long as both jobs use the same encryption password and back up VMs to the same vault, moving a VM from one job to another does not cause the VM to reseed.

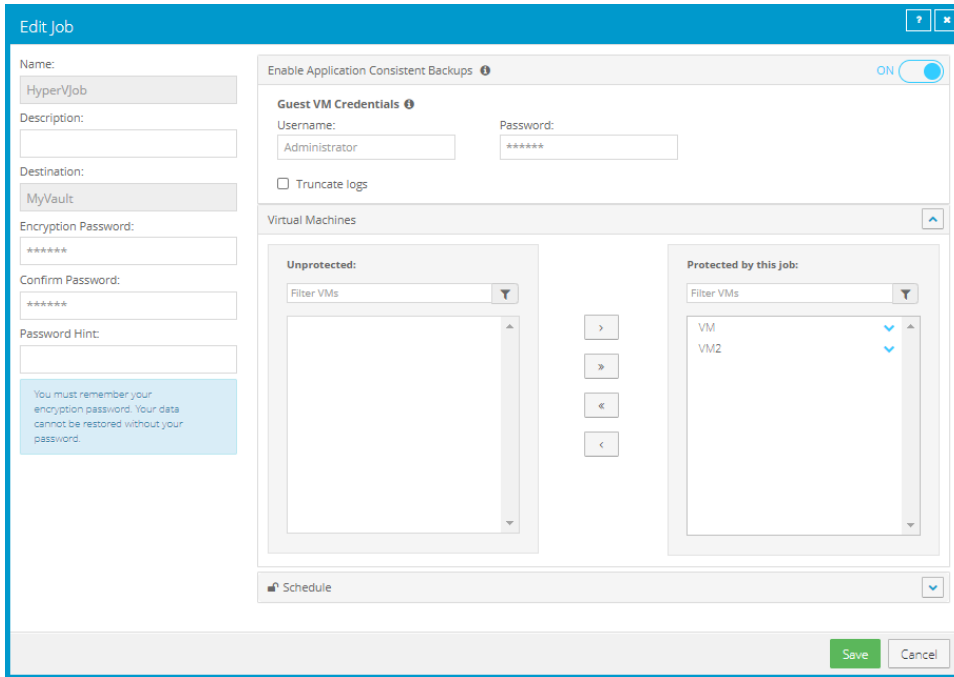
To edit a Hyper-V backup job:

1. On the navigation bar, click **Computers**.

The Computers page shows registered computers and environments.


2. Click the Hyper-V environment row. 
3. Do one of the following:
 - On the **Jobs** tab, in the **Jobs** column, click the name of the job that you want to edit.
 - On the **Jobs** tab, find the job that you want to edit. In its **Select Action** menu, click **Edit Job**.
 - On the **Virtual Machines** tab, in the **Jobs** column, click the name of the job that you want to edit.
 - On the **Virtual Machines** tab, click a VM that belongs to the job you want to edit. In the **Select Action** menu, click **Edit Job**.
4. If required, change the job description, encryption password or password hint at the left side of the Edit Job dialog box.
5. To change credentials for application-consistent backups on VMs in the job, do one or both of the following:
 - To change credentials for VMs in the job, enter an admin user's username and password in the Guest VM Credentials area.
 - To change credentials for a specific VM in the job, click the arrow at the right side of the VM name in the **Protected by this job** box, and enter an admin user's username and password for the VM.


The specified user must have admin access to VMs in the backup job. You can enter the username as *username* or *domain\username*.



6. Do one or more of the following to add VMs to or remove VMs from the job:

- To find one or more VMs in the **Unprotected** or **Protected by this job** box, enter characters from the VM names in the associated **Filter VMs** box.


- To add all VMs in the **Unprotected** box to the backup job, click **Protect all.** 

- To add some VMs in the **Unprotected** box to the backup job, select the VMs in the **Unprotected** box, and then click **Protect selected.** 

To select multiple VMs in the list, press CTRL and click the VM names. To select multiple consecutive VMs in the list, press Shift and then click the first and last VM that you want to select, or drag the mouse across the VMs.

- To remove all VMs in the **Protected by this job** box from the backup job, click **Unprotect all.**



- To remove some VMs in the **Protected by this job** box from the backup job, select the VMs in the **Protected by this job** box, and then click **Unprotect selected.** 

To select multiple VMs in the list, click each VM name.

Ensure that each VM that you want to include in the backup job appears in the **Protected by this job** box.

Note: When editing a Hyper-V job, you cannot select VMs from a list to include or exclude, as you can when adding a job.

7. To schedule the backup job to run, click **Schedule**. In the **Schedule** box that appears, create one or more schedules. See [Add or edit a schedule for a Hyper-V backup job](#).
8. Click **Save**.

5.6 Add or edit a schedule for a Hyper-V backup job

When adding or editing a Hyper-V backup job, you can create a schedule for running the job, and enable or disable the schedule. You can also edit existing schedules.

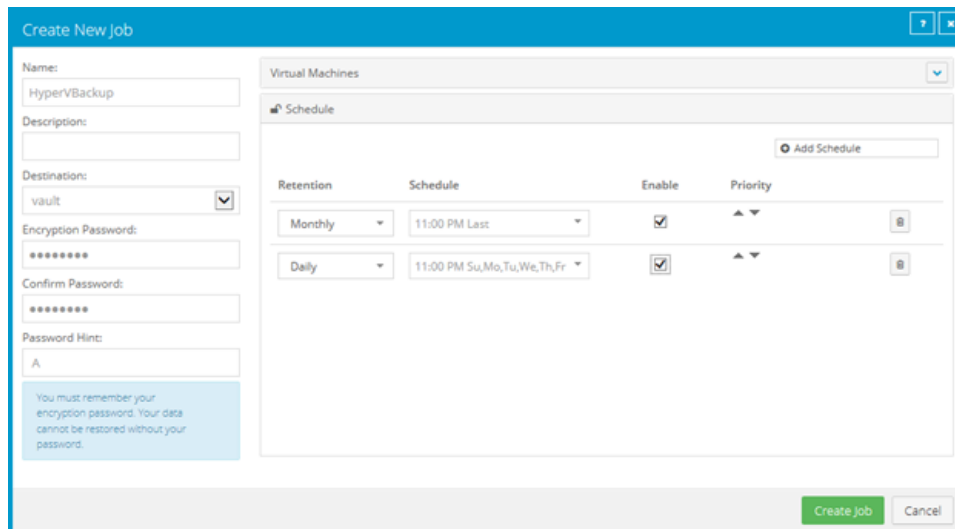
You can specify a retention type for each schedule. The retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline. See [Add retention types](#).

You can create complex schedules for a Hyper-V backup job by creating multiple schedules. For example, you can schedule a backup job to run at midnight every Friday, and schedule the job to run at 8 pm on the first day of every month. To create multiple schedules, repeat the following procedure for each schedule.

If a job is scheduled to start at exactly the same time by multiple schedules, the job only runs once at the scheduled time. If the jobs have different retention types, the retention type of the schedule that is highest in the list is applied to the resulting safeset. For example, in the following screenshot, the job is scheduled to run at 11 PM on the last day of the month with the Monthly retention type, and every night at 11 PM with the Daily retention type. On the last day of each month, the job runs only once at 11 PM. Because the schedule with the Monthly retention type is higher in the list than the schedule with the Daily retention type, the Monthly retention type is applied to the safeset.

Note: If a Hyper-V VM is being backed up when a second backup starts for the same VM, the second backup fails; it is not queued. For example, if a VM is scheduled to be backed up at 11 PM by one schedule and at 11:01 PM by another schedule, the second backup attempt fails. Try to avoid overlapping schedules; problems can occur if a job is scheduled to run twice in a short period of time.

To run a backup at any time, without scheduling it, see [Run an ad-hoc backup](#).



Note: You cannot defer scheduled Hyper-V backups. Hyper-V Agent backups can only be deferred when they are run manually (ad hoc).

To add or edit a schedule for a Hyper-V backup job:

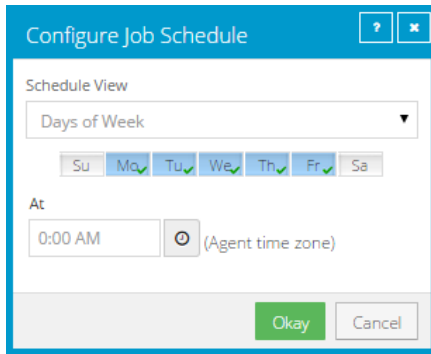
1. In the Create New Job or Edit Job dialog box, while adding or editing a Hyper-V backup job, click **Schedule**.
2. In the **Schedule** box, do one of the following:
 - To add a schedule, click **Add Schedule**.
 - To edit a schedule, find the schedule that you want to edit.
3. In the schedule row, select a retention type.

A retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.


4. In the **Schedule** box, click the arrow.

The Configure Job Schedule dialog box opens.


5. In the Configure Job Schedule dialog box, do one of the following:
 - To run the backup on specific days each week, click **Days of Week** in the **Schedule View** list. Select the days when you want to run the job. Use the **At** field to specify the time when you want to run the job each day. Click **Okay**.




- To run the backup on specific dates each month, click **Days of Month** in the **Schedule View** list. On the calendar, select the dates when you want to run the job. Use the **At** field to specify the time when you want to run the job on each date. Click **Okay**.



- To create a custom schedule, click **Custom** in the **Schedule View** list. In the **Custom Cycle** box, enter a custom schedule. Follow the format and notation described in the dialog box. Click **Okay**.



The new or revised schedule appears in the **Schedule** box.

6. To enable the schedule to run, select **Enable**. To disable the schedule so it does not run, clear **Enable**.
7. To remove the schedule, click **Delete Schedule**. 
8. If there is more than one schedule row, you can use the **Priority** arrows to move a schedule higher or lower in the list. Schedules that are higher in the list have higher priority than schedules lower in the list. If a job is scheduled to run at the same time by multiple schedules, the job only runs

once at the scheduled time. If the schedules have different retention types, the job only runs with the retention type of the schedule that is highest in the list.

9. Click **Save**.

5.7 Disable or enable all scheduled backup jobs

Admin users can disable or enable all scheduled backup jobs for a protected environment.

Note: You can also disable or enable a specific schedule for a backup job by clearing or selecting the schedule's **Enable** check box. See [Add or edit a schedule for a Hyper-V backup job](#).

When you disable all scheduled jobs for a computer or protected environment, backup jobs do not run according to any schedules. When jobs are disabled for most computers and protected environments, the **Enable** check box in the View/Add Schedule dialog box is cleared. When jobs are disabled for a Hyper-V environment, you cannot view or edit schedules in the **Schedule** area of the Edit Job dialog box.


When you enable scheduled jobs for most computers or protected environments, the **Enable** check box in the View/Add Schedule dialog box is selected for all schedules, and all jobs run according to all schedules.

Enabling all scheduled jobs can be particularly useful after a Hyper-V disaster recovery. When you recover jobs and settings from an offline Hyper-V Agent, all scheduled backup jobs for the Agent are disabled.

When you enable scheduled jobs for a Hyper-V environment, jobs run according to any schedules where the **Enable** check box is selected in the Edit Job dialog box.

To enable or disable all schedules:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Select the check box to the left of each computer or protected environment for which you want to enable or disable all schedules.



3. In the **Actions** list, do one of the following:
 - To enable all schedules for the selected computers, click **Enable Scheduled Jobs**.
 - To disable all schedules for the selected computers, click **Disable Scheduled Jobs**.

5.8 Run an ad-hoc backup

After a backup job is created, you can run the backup at any time, even if the job is scheduled to run at specific times.

To run an ad-hoc backup:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the computer with the backup job that you want to run, and expand its view by clicking the computer row.
3. Click the **Jobs** tab.
4. Find the job that you want to run, and click **Run Job** in its **Select Action** menu.

The Run Job dialog box shows the default settings for the backup.

Note: Beginning at this point, you can click **Start Backup** to immediately start the job. If you prefer, you can change backup options before running the job.

5. To back up the data to the vault specified in the job, do not change the **Destination**.

Note: You cannot change the destination for Hyper-V Agent backups. SSI files are not supported for Hyper-V jobs.

6. In the **Retention Scheme** list, click a retention type.

The retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.

7. Do one of the following:
 - To allow the backup job to run without a time limit, clear the **Use Deferring** check box.
 - To specify a maximum amount of time that the backup job can run, select the **Use Deferring** check box. From the **Backup time window** list, select **Minutes** or **Hours**. In the adjacent box, type the maximum number of minutes or hours that the job can run.

Note: When deferring is used, the backup job does not back up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the backup time window.

8. Click **Start Backup**.

The Process Details dialog box shows the backup progress, and indicates when the backup is completed. Other recent job processes might also be listed in the dialog box. See [View current process information for a job](#).

9. If you want to stop the backup, click **Stop**.

10. To close the Process Details dialog box, click **Close**.

6 Restore Hyper-V data

When VMs are protected in a Hyper-V environment, you can:

- [Restore Hyper-V VMs](#)
- [Restore a Hyper-V VM within minutes using Rapid VM Restore](#)

Hyper-V Agent 9.00 or later is required for Rapid VM Restores.

- [Restore Hyper-V files, folders and database items](#)

Hyper-V Agent version 8.84 or later is required for restoring specific files, folders and database items from protected Hyper-V VMs.

6.1 Restore Hyper-V VMs

You can restore one or more virtual machines (VMs) from a Hyper-V backup. In a single request, you can restore up to 50 VMs— even if they were backed up using multiple Hyper-V backup jobs with different encryption passwords.

When restoring a VM, you must specify a destination for the VM files. If you are restoring to a Hyper-V cluster, available destinations are Cluster Shared Volumes (CSV) found in the Failover Cluster Manager. If you are restoring to a standalone host, available destinations are volumes on direct attached storage.

You can also specify a datastore folder for the files. If you do not specify a folder, a new folder with the same name as the VM is created for the VM files. All of a VM's disks are restored in a single location, even if the disks originally resided on different volumes and you select the original host and datastore. Hyper-V VM files will not be restored to a non-empty folder. If a folder exists, it will create a new folder. If you force a custom folder, and this folder exists, the restore of that specific VM will fail.

You can only restore a Hyper-V VM to a host where the Host service is installed. When restoring a Hyper-V VM in a cluster, you must choose a host where the Host service is running or the restore will fail. If the Host service is not installed on the node where you want to restore a VM, you can restore the VM to a node that has the Host service, and then migrate the VM to another node in the cluster.

Note: Portal does not indicate which hosts in a cluster have the Host service installed. All hosts in a Hyper-V cluster appear on the Hosts tab on the Computer page, even if the Host service is only installed on some of the hosts. If the status of a host is "Offline", the Host service is either not installed or not running on the host. We recommend installing the Host service on each host in a cluster. See [Recommended deployment for protecting a Hyper-V cluster](#).

Restored VMs are imported automatically into Hyper-V. Restored VMs keep their original names, unless you specify new VM names during the restore process.

Each VM has a unique identifier. You can restore a VM with its original internal identification number (GUID), with a new GUID, or with a new GUID if a VM with the original GUID exists in the Hyper-V environment.

Note: A restored Hyper-V VM never overwrites an existing VM.

Beginning in version 8.84 of the Hyper-V Agent, if a VM was backed up with an ISO image file connected to its DVD drive, the ISO image file is not connected to the DVD drive when the VM is restored.

The generation of a VM is retained when it is backed up and restored. A protected Generation 1 VM is restored as a Generation 1 VM. A protected Generation 2 VM is restored as a Generation 2 VM.

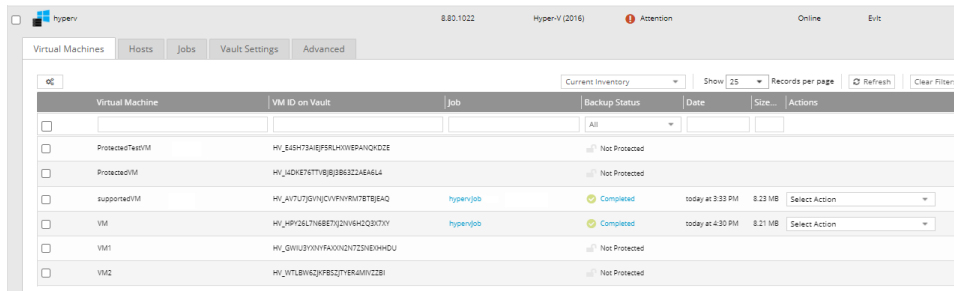
If you stop a restore process, VMs that are restored before you stop the process remain in the Hyper-V environment. VMs that are not fully restored when you stop the process are not restored.

Note: If you stop a restore and one or more VMs are not completely restored, partial VM files will remain on disk. You must clean up these files manually.

To restore Hyper-V VMs:

1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the Hyper-V environment with the VM that you want to restore, and expand the environment view by clicking the row.
3. Click the Virtual Machines tab.

The Virtual Machines tab shows VMs in the Hyper-V environment.



Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Size	Actions
ProtectedTestVM	HY_E4B473A8F5RLK5WEPANQDZE		Not Protected			
ProtectedVM	HY_H4D4E767YB8J32246464		Not Protected			
supportedVM	HY_A1V7JGVWQVFNRYN4T8TBEAQ	hyperVJob	Completed	today at 3:33 PM	8.23 MB	Select Action
VM	HY_HRYZBL7N6BET7ZNV6HQZB7YX	hyperVJob	Completed	today at 4:30 PM	8.21 MB	Select Action
VM1	HY_GWVUB3YXVFAOXZNTZ5NEHMDU		Not Protected			
VM2	HY_V7L8V6ZKPB8QTYER8M1Z2BI		Not Protected			

4. In the Current Inventory/Protected Inventory filter, click **Protected Inventory**.

The Virtual Machines tab shows VMs that have been backed up and can be restored.

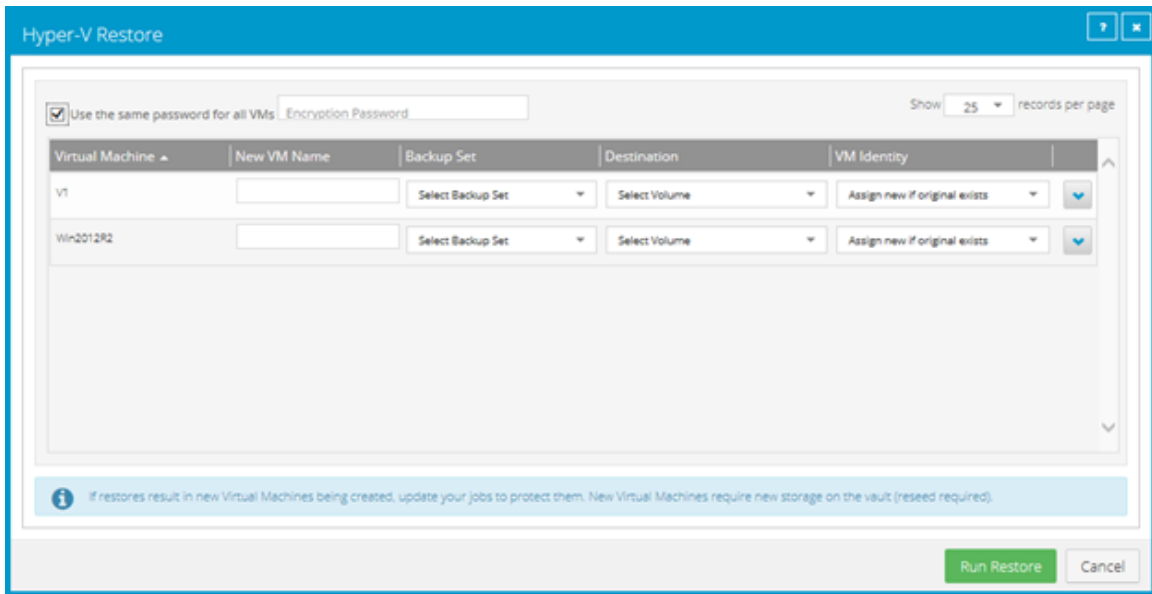
5. Do one of the following:
 - To restore one VM, click **Restore** in its **Select Action** menu.
 - To restore multiple VMs, select the check box for each VM that you want to restore. Click **Restore Hyper-V Job**.

Restore Hyper-V Job. 

You can restore up to 50 VMs at a time.


6. In the Choose What You Want to Restore dialog box, select **Virtual Machines**, and then click **Continue**.

The Hyper-V Restore dialog box shows the VM or VMs that you want to restore.



7. Do one of the following:

- If you are restoring one VM, go to the next step.
- If you are restoring multiple VMs protected with the same encryption password, select the **Use the same password for all VMs** check box. In the **Encryption Password** box, enter the data encryption password.

To view a password hint, click the **Hint** button. 

- If you are restoring multiple VMs that were protected by jobs with different encryption passwords, clear the **Use the same password for all VMs** check box.

8. In the row for each VM that you are restoring:

- (Optional) In the **New VM Name** box, enter a name for the restored VM. If you do not enter a name, the VM is restored with its original name.
- In the **Backup Set** list, click the backup from which you want to restore. If you did not enter the same password for all VMs, enter the password in the **Encryption Password** box. Click **Apply**.
- In the **Destination** list, click the destination for the VM files. If you want to specify a folder for restoring the VM files, enter the folder in the **Sub-Path** box. Click **Apply**.

In a Hyper-V cluster, you can restore files to a CSV. In a standalone host, you can restore files to volumes on direct attached storage. You cannot restore VMs to volumes smaller than 5 GB in size, or to system volumes. These volumes do not appear in the Destination list.

If you do not have sufficient space for a restore, you can add additional storage. Newly-added storage should be available in the Destination list immediately. However, if you do not see a new storage device in the Destination list, stop and restart the Hyper-V services.

If you do not specify a folder, the VM is restored to a folder with the VM name.

You can also enter subfolders in the **Sub-Path** box (e.g., folder\subfolder1\subfolder2).

Note: Hyper-V VM files will not be restored to a non-empty folder. If a folder exists, it will create a new folder with the same name followed by a number in brackets ().

- In the **VM Identity** list, do one of the following:
 - To restore the VM with a new GUID, click **Assign new identity**.
Note: If a node is down but has not been evicted from the cluster, you can only restore the VM using the **Assign new identity** option. This prevents the VM from being restored with the same GUID as a VM on the cluster node that is down.
 - To restore the VM with its original GUID, click **Restore original identity**.
Note: If a VM with the original GUID exists in the Hyper-V environment, the restored VM will not overwrite the existing VM. Two VMs in a Hyper-V environment can have the same GUID if they are on separate hosts and are not configured for high availability.
 - To restore the VM with its original GUID unless a VM with the original GUID exists in the Hyper-V environment, click **Assign new if original exists**. If a VM with the original GUID exists in the Hyper-V environment, the VM is restored with a new GUID.
- If the **Host** list is not shown, click the VM row to expand its view. Do one or more of the following:
 - To specify a host for the restored VM, click a host in the **Host** list.
 - To power on the VM after it is restored, select **Power on VM**.
 - To leave the restored VM powered off, clear **Power on VM**.
 - To connect the restored VM to the network, select **Enable network connectivity**.
If **Enable network connectivity** is selected, and the VM has a network adapter with the same name as a network adapter on the host, the VM will be automatically connected to the network.
 - To restore the VM without network connectivity, clear **Enable network connectivity**.

9. Click **Run Restore**.

6.2 Restore a Hyper-V VM within minutes using Rapid VM Restore

Using Rapid VM Restore, you can restore a virtual machine (VM) to your Hyper-V environment within minutes. You can only restore one VM at a time using this restore method.

When you first restore a Hyper-V VM using Rapid VM Restore, disks from the protected VM are mounted on a temporary VM for immediate access. While the VM runs, changes are written to a temporary storage location. At this stage, the VM requires a running Rapid VM Restore process and connections to the Hyper-V Agent and vault, and is intended for temporary use.

You can restore the VM permanently by migrating it to a permanent storage location using Portal. After a VM is migrated, the VM does not require a running Rapid VM Restore process and is independent from the Hyper-V Agent and vault. See [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).

We recommend migrating a VM to permanent storage as soon as possible after it is restored using Rapid VM Restore. If the network connection to the Hyper-V Agent, vault or destination host is interrupted before a VM is migrated to permanent storage, VM data could be lost.

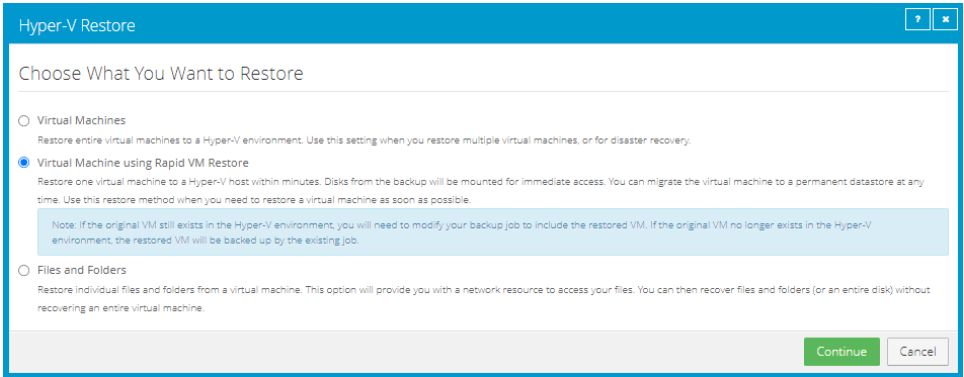
Notes:

- A Hyper-V VM restored using Rapid VM Restore cannot be backed up until it is migrated to permanent storage. If you try to back up the restored VM before it is migrated to permanent storage, the following error occurs: *Unable to backup virtual machine "VMname" [VMID] because it is in RVMR state*. The VM's backup status in Portal is *Failed*.
- When you first restore a VM using Rapid VM Restore, it runs on the host that you select during the restore process. The VM could be lost if it is migrated to another host in a Hyper-V cluster. For this reason:
 - High availability is not enabled for a VM that is running using Rapid VM Restore. If high availability was enabled for the VM when it was backed up, high availability will be enabled for the VM after it is migrated to permanent storage using Portal. See [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).
 - Do not enable high availability on a VM that is running using Rapid VM Restore.
 - Only use Portal to migrate a VM that is running using Rapid VM Restore. Do not migrate the VM to another host in a cluster using Hyper-V Manager.
- Rapid VM Restore is available with Hyper-V Agent version 9.0 or later.

To restore a Hyper-V VM within minutes using Rapid VM Restore:

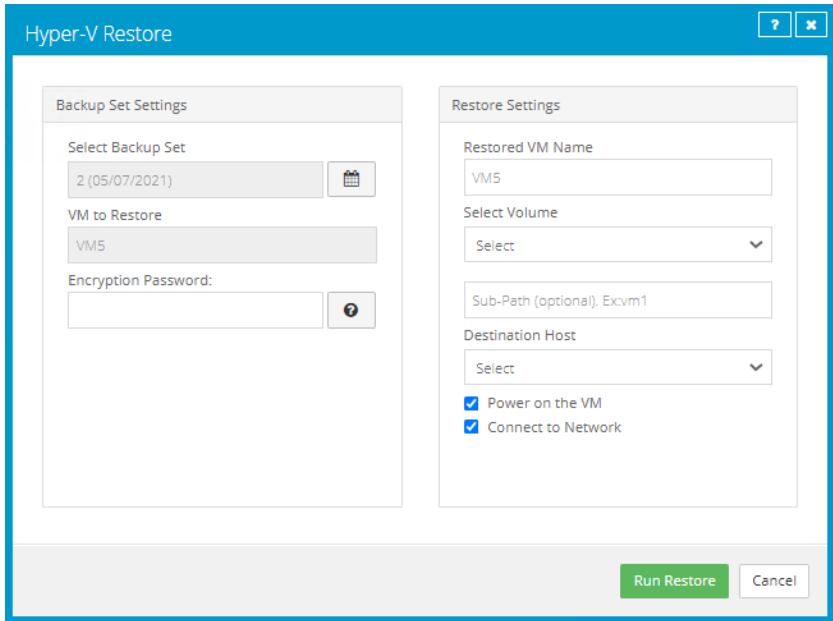
1. On the navigation bar, click **Computers**.
A grid lists available computers.
2. Find the Hyper-V environment with the VM that you want to restore, and expand the environment view by clicking the row.
3. Click the Virtual Machines tab.
The Virtual Machines tab shows all VMs in the Hyper-V environment.
4. In the Current Inventory/Protected Inventory filter, click **Protected Inventory**.
The Virtual Machines tab shows VMs that have been backed up.
5. Find the VM that you want to restore, and click **Restore** in the job's **Select Action** menu.
6. In the Choose What You Want to Restore dialog box, select **Virtual Machine using Rapid VM Restore**.


If the **Virtual Machine using Rapid VM Restore** option does not appear, this restore method is not available. This could occur with a Hyper-V Agent version earlier than 9.00, with a Portal version earlier than 8.89, or if backups are not available in a local vault that supports Rapid VM Restores. For complete requirements, see [Hyper-V Rapid VM Restore requirements](#).




7. Click **Continue**.

The Hyper-V Restore dialog box appears. The Select Backup Set box shows the most recent safeset for the VM on a vault that supports Rapid VM Restores. The VM to Restore box shows the VM that you are restoring.



8. To restore from an older safeset, click the **Browse Safesets** button.  In the calendar that appears, click the date of the safeset from which you want to restore. In the safeset list to the right of the calendar, click the safeset from which you want to restore. The list only includes safesets on vaults that support Rapid VM Restores.

9. In the **Encryption Password** box, enter the data encryption password. To view the password hint, click the **Hint** button. 
10. In the Restore Settings box, do the following:
 - In the **Restored VM Name** box, type a name for the restored VM.

If you specify the name of a VM that already exists in the Hyper-V environment (e.g., the VM that was backed up), the restored VM will have the following name: *VMname-rvmr-yyyy-Mon-dd--hh-mm-ss*, where *yyyy-Mon-dd--hh-mm-ss* is the date and time when the VM was restored (e.g., VM-rvmr-2019-Nov-27--06-14-09).
 - In the **Select Volume** list, select a volume for writing changes while the VM is running using Rapid VM Restore but is not migrated to permanent storage. The amount of free space is shown for each volume in the list.
 - (Optional) In the Sub-Path box, type the folder path (e.g., *RestoredVMs\VM 1*) on the selected volume for writing changes while the VM is running using Rapid VM Restore.


If you do not specify a path, changes will be written to a folder with the name of the restored VM: *VMname-rvmr-yyyy-Mon-dd--hh-mm-ss*, where *yyyy-Mon-dd--hh-mm-ss* is the date and time when the VM was restored.
 - In the **Destination Host** list, select a host for running the restored VM.
 - Do one of the following:
 - To restore the VM with its power on, select the **Power on the VM** option.
 - To restore the VM powered off, clear the **Power on the VM** option. Restoring a VM with the power off can be useful if you want to verify or change the VM settings before powering on the VM.

Note: If you are restoring a VM that still exists in the Hyper-V environment, power off the original VM before the restore to avoid conflicts between the original VM and the restored VM.
 - Do one of the following:
 - To connect the VM to the network, select **Connect to Network**.
 - To restore the VM without network connectivity, clear **Connect to Network**. Restoring a VM with the power off can be useful when restoring to a Hyper-V environment that does not have the original network. You can then verify the VM settings before connecting the VM to the network.
11. Click **Run Restore**.

The Process Details dialog box appears. When the VM is restored, the following Status message appears in the Process Details dialog box: *Rapid VM restore is running*.


Note: Record the Process ID of the restore. If the same VM is restored more than once concurrently using Rapid VM Restore, you can use the Process ID to identify the restored VM.

The restored VM appears in the Hyper-V environment. You can access the VM and begin using it.



The restored VM also appears in the list of unprotected VMs on the Computers page in Portal. You can add the VM to a Hyper-V backup job but you cannot back it up until it has been migrated to permanent storage. See [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).

12. Do one or more of the following:

- To close the Process Details dialog box, click **Close** in the dialog box. If you close the Process Details dialog box without canceling the Rapid VM Restore, the VM remains in the Hyper-V environment.
- To reopen the Process Details dialog box, find the VM you are restoring on the Virtual Machines tab of the Hyper-V environment on the Computers page. Click the Rapid VM Restore symbol that appears beside the VM name: 
- To permanently restore the VM by migrating it to permanent storage, see [Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage](#).
- To remove the VM from the Hyper-V environment, click **Cancel Rapid VM Restore** in the Process Details dialog box.

6.2.1 Migrate a Hyper-V VM restored using Rapid VM Restore to permanent storage

When you first use Rapid VM Restore to restore a Hyper-V VM, the VM is dependent on the Hyper-V Agent and vault, and is intended for temporary use.

To restore the VM permanently, use Portal to migrate the VM to a permanent storage location in the Hyper-V environment. You can migrate the VM to a location on the same volume selected for the Rapid VM Restore or to a different volume.


If the VM is powered on, you can continue to use the VM during the migration. After migration, the VM is independent from the Hyper-V Agent and vault.

If you cancel a migration before a VM is fully migrated to the permanent location, the restored VM remains in the Hyper-V environment and continues running using the Rapid VM Restore process. If you do not cancel the Rapid VM Restore process, you can try to migrate the VM again.

Notes:

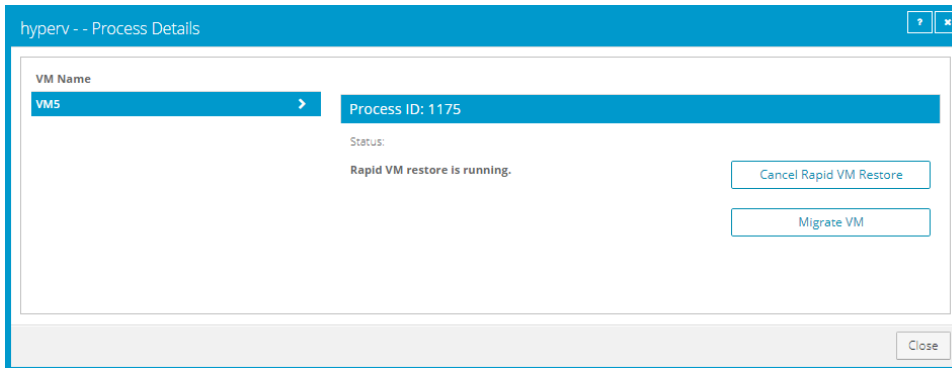
- We recommend using Portal to migrate VMs to permanent storage rather than using Hyper-V Manager. If you restore a VM using Rapid VM Restore and migrate it to a different host and storage using Hyper-V Manager, you will not be able to migrate the VM to permanent storage using Portal.
- A Hyper-V VM restored using Rapid VM Restore cannot be backed up until it is migrated to permanent storage.
- While a VM is being migrated, you cannot power on, power off, suspend or create a checkpoint for the VM using the Hyper-V Manager.
- If high availability was enabled for a VM when it was backed up, high availability will be enabled for the restored VM after it is migrated to permanent storage. However, specific high availability settings (e.g., preferred owner, failover and failback settings) are not applied.

To migrate a Hyper-V VM restored Rapid VM Restore to permanent storage:

1. If the Process Details dialog box is not open for the Rapid VM Restore process, find the Hyper-V environment on the Computers page. On the Virtual Machines tab for the Hyper-V environment, click the Rapid VM Restore symbol beside the protected VM name: 

Note: The Rapid VM Restore symbol appears beside the VM that was backed up, not beside the VM restored using Rapid VM Restore.

The Process Details dialog box shows Rapid VM Restore processes for the VM. If the VM is restored more than once concurrently using Rapid VM Restore, the protected VM name appears more than once in the VM Name list at the left of the dialog box.

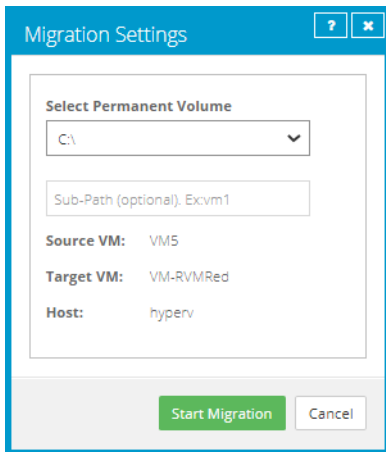


2. If the protected VM name appears more than once in the VM Name list, check that the process ID for the VM that you want to migrate (recorded in [Restore a Hyper-V VM within minutes using](#)

[Rapid VM Restore](#)) is shown in the middle of the dialog box. If the correct process ID is not shown, click another VM name in the VM Name list.

3. Click **Migrate VM**.

The Migration Settings dialog box appears. If you specified a Sub-Path when starting the Rapid VM Restore, this location is populated in the dialog box.



4. In the **Select Permanent Volume** list, select the permanent storage volume for the VM files.
5. (Optional) In the Sub-Path box, type the folder path (e.g., `RestoredVMs\VM 1`) on the selected volume for permanently storing VM files. If you do not specify a path, files will be saved in a folder with the restored VM's name on the selected volume.
6. Click **Start Migration**.

The following Status message appears in the Process Details dialog box: *VM migration is in progress*.

If you click **Cancel Migration** while the migration is in progress, the restored VM remains in the Hyper-V environment and is still dependent on the Hyper-V agent and vault. You can start the migration again, if desired.

When the VM is migrated to the permanent storage location, the following Status message appears in the Process Details dialog box: *VM has been migrated*. At this point, the VM is permanently restored and is no longer dependent on the Hyper-V agent and vault. The Rapid VM Restore process ends and the Rapid VM Restore symbol no longer appears beside the protected VM name on the Computers page.

6.3 Restore Hyper-V files, folders and database items

Beginning with version 8.84 of the Hyper-V Agent, you can restore files and folders from protected Windows VMs in Hyper-V environments.

During a file and folder restore, volumes from a protected VM are mounted in a RestoreMount folder on the server where the Management service is running. The folder is shared, and a UNC path to the share is

provided in Portal. You can then access the share from a VM or server with network access to the server, and copy files and folders that you want to restore from the protected VM.

To restore files and folders from Windows VMs, the Hyper-V Agent Management service must be installed on the same Windows version or a later version than is installed on the Windows VMs. For example, to restore files and folders from Windows Server 2019 VMs, the Hyper-V Agent Management service must be installed on Windows Server 2019 or later.

You can select multiple VMs in a single file and folder restore. When you restore files and folders from multiple VMs, a separate UNC path is provided in Portal for each VM.

Note: To access a UNC share during a file and folder restore, you must provide credentials for a user with admin access to the server where the Management service is running.

In addition to copying files and folders from a protected VM, you can find and restore items from Exchange and SQL Server databases. Using the Granular Restore for Microsoft Exchange and SQL application, you can restore Exchange mailboxes and messages to PST files or live databases, export SQL Server database items to live databases, and export SQL Server database items as SQL scripts. For more information, see the *Granular Restore for Microsoft Exchange and SQL User Guide*.

Note: You cannot restore files, folders and database items from Linux VMs in Hyper-V environments.

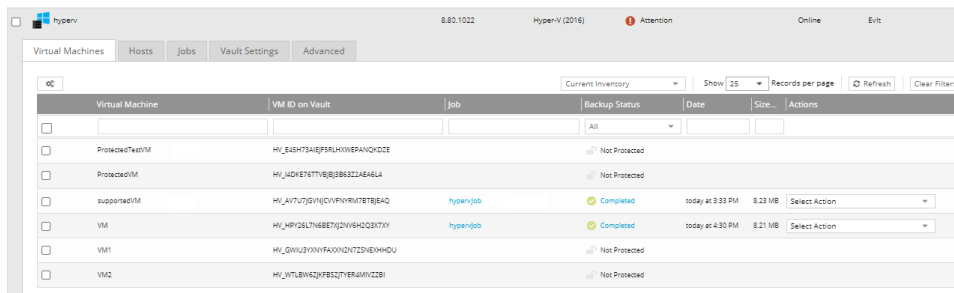
To restore Hyper-V files, folders and database items:

1. On the navigation bar, click **Computers**.

A grid lists available computers.

2. Find the Hyper-V environment with the VM that you want to restore, and expand the environment view by clicking the row.
3. Click the Virtual Machines tab.


The Virtual Machines tab shows all VMs in the Hyper-V environment.



4. In the Current Inventory/Protected Inventory filter, click **Protected Inventory**.

The Virtual Machines tab shows VMs that have been backed up.

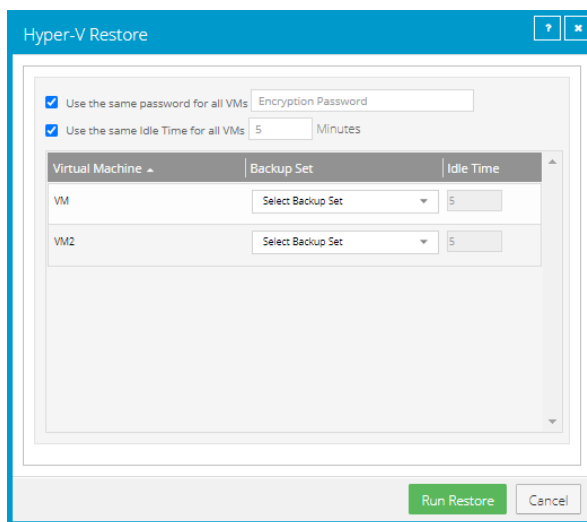
5. Do one of the following:

- To restore files and folders from one VM, click **Restore** in its **Select Action** menu.
- To restore files and folders from one or more VMs, select the check box for each VM, and then click **Restore Hyper-V Job**. 

6. In the Choose What You Want to Restore dialog box, select **Files and Folders**.

The Hyper-V Restore dialog box shows the VM or VMs from which you want to restore files and folders. If you are restoring files and folders from multiple VMs, encryption password and idle time options appear at the top of the dialog box.

If you are restoring files and folders from one VM, go to [Step 9](#).



7. If you are restoring files and folders from multiple VMs, do one of the following:

- If the VMs are protected with the same encryption password, select the **Use the same password for all VMs** check box. In the **Encryption Password** box, enter the data encryption password.
- If the VMs are protected by jobs with different encryption passwords, clear the **Use the same password for all VMs** check box.

8. If you are restoring files and folders from multiple VMs, do one of the following:

- To set the same idle time for each VM, select the **Use the same Idle Time for all VMs** check box. In the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will automatically unshare. The idle time value can be from 2 to 180 minutes.

Note: The drive will not unshare as long as new data is being copied. If you copy the same data from a shared drive more than once, the system could time out because no new data is being read.

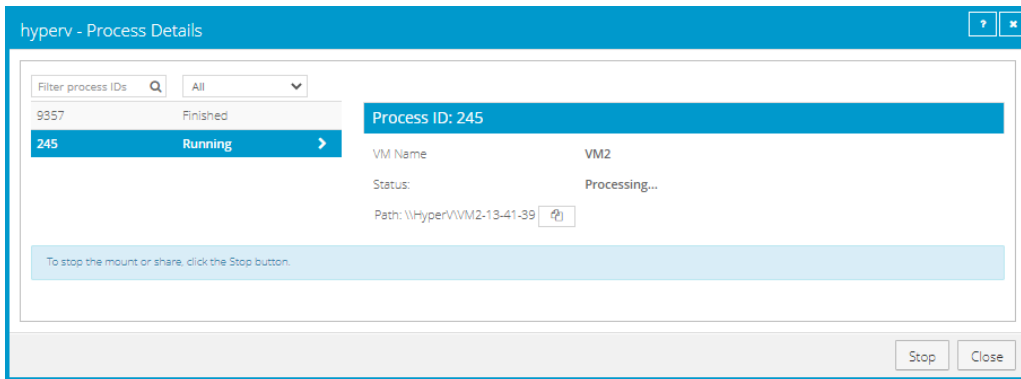
- To set a different idle time for each VM, clear the **Use the same Idle Time for all VMs** check box. You can then set the idle time for each VM in [Step 9](#).
9. For each VM from which you are restoring files and folders, do the following in the VM row:
- In the **Backup Set** list, click the backup from which you want to restore. If you did not enter an encryption password for all VMs in [Step 7](#), enter the password in the **Encryption Password** box. Click **Apply**.
 - If you did not specify an idle time for all VMs in [Step 8](#), in the **Idle Time** box, enter the number of minutes of inactivity after which the shared drive will automatically unshare. The idle time value can be from 2 to 180 minutes.


Note: The drive will not unshare as long as new data is being copied. If you copy the same data from a shared drive more than once, the system could time out because no new data is being read.

10. Click **Run Restore**.

The Process Details dialog box shows the process status. If you are restoring files and folders from multiple VMs, a separate process appears for each VM. To view the process status for another VM, click the running process on the left side of the Process Details dialog box.

When VM volumes are shared, a UNC path to the share appears in the dialog box. The path is named `//hostName/vmName-hh-mm-ss`, where `hh-mm-ss` is the time when the share was created on the server where the Management service is running.



11. To copy the UNC path, click the Copy Path to Clipboard button .

If you are restoring files and folders from multiple VMs, a different UNC path is provided for each VM. To obtain the UNC path for another VM, click the running process on the left side of the Process Details dialog box.

12. Use the UNC path on a VM or server with network access to the server where the Management service is running to do one or both of the following:
 - Access volumes from the protected VM, and copy files and folders that you want to restore.
IMPORTANT: To access the UNC share, you must provide credentials for a user with admin access to the server where the Management service is running.
 - Use the Granular Restore for Microsoft Exchange and SQL application to find and restore items from Exchange and SQL Server database backups in the mounted volumes. You can restore Exchange mailboxes and messages to PST files or live databases, export SQL Server database items to live databases, and export SQL Server database items as SQL scripts. See the *Granular Restore for Microsoft Exchange and SQL User Guide*.

7 Recover jobs and settings from an offline Hyper-V Agent

You can recover jobs and settings from an offline Hyper-V Agent:

- During a disaster recovery.
- When moving to a new Hyper-V environment.
- When a Hyper-V agent is not connecting to Portal because of a Portal certificate change. If a Hyper-V agent is not connecting to Portal and a *The Agent Management SSL Certificate does not match what was expected* message appears in the Host log, please contact your service provider or Portal administrator to determine whether you need to recover the agent's jobs and settings. If this is required, you can back up Hyper-V agent logs in the <ManagementServiceInstallFolder>\Data folder, uninstall the Hyper-V agent that is not connecting to Portal, and then recover the agent's jobs and settings.

You can install a new Hyper-V agent and recover the following information and settings from an offline Hyper-V agent:

- Backup jobs
- Vault settings
- Hyper-V environment address and last backup status
- Advanced settings, including the Agent description, retention types, notifications, and bandwidth throttling

You can then enter credentials, run backup jobs from the original agent, and restore VMs that were protected by the agent.

You cannot recover passwords for a Hyper-V Agent. You must manually enter Hyper-V environment, vault, and encryption passwords after recovering Hyper-V Agent jobs and settings. You might also need to enter passwords for application-consistent backups and an SMTP password for notifications.

IMPORTANT: You must enter these passwords when recovering a Hyper-V Agent even though asterisks appear in the Portal password fields.

To recover jobs and settings from an offline Hyper-V Agent:

- the newly-installed Hyper-V Agent must be the same version or a later version than the offline agent. For example, you can install a version 8.84 agent and recover jobs and settings from an offline version 8.80 agent.
- the new Hyper-V environment must be the same version or a later version than the environment protected by the offline agent. For example, you can recover jobs and settings from an offline Hyper-V agent in a Windows 2012 R2 environment to an agent in a Windows Server 2016 environment.

The generation of a VM is retained when it is backed up and restored. A protected Generation 1 VM is restored as a Generation 1 VM. A protected Generation 2 VM is restored as a Generation 2 VM.

When you recover jobs and settings from an offline Hyper-V Agent, all scheduled backup jobs for the Agent are disabled. If Hyper-V VMs remain in the protected environment, or have been restored after a disaster, you can re-enable all scheduled jobs for the environment. See [Disable or enable all scheduled backup jobs](#).

IMPORTANT: Hyper-V Agent settings are saved in the Portal database. To ensure that a Hyper-V environment can be fully restored if the Portal is also lost, the Portal database must be backed up. For more information, see the *Portal Installation and Configuration Guide*.

To recover jobs and settings from an offline Hyper-V Agent:

1. Install the Hyper-V Agent Management service on a supported Windows server. See [Install the Hyper-V Agent Management service](#).

On the Register Hyper-V Agent Management with Portal page of the installer, register the Management service to the Portal where the original Hyper-V Agent was registered. Register the Management service to the Portal using the user who installed the original Hyper-V Agent, or using an Admin user in the original user's site.

2. Log in to Portal as the user who installed the original Hyper-V Agent, or as an Admin user in the original user's site.

3. In Portal, on the navigation bar, click **Computers**.

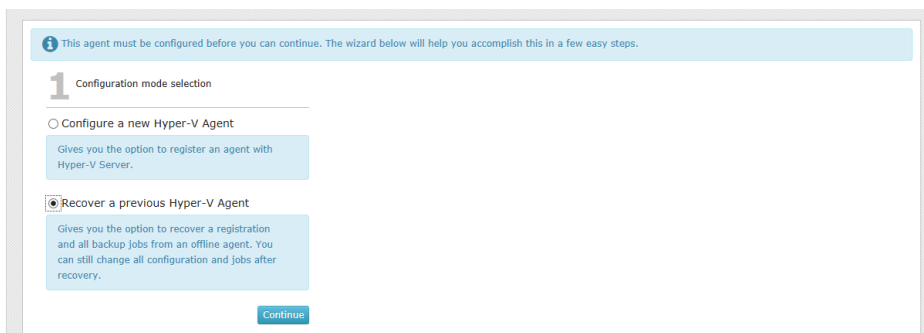
The Computers page shows registered computers.

4. Find the computer where the new Hyper-V Agent Management service is installed, and expand its view by clicking its row.

Before you recover jobs and settings from the offline Hyper-V agent, the name of the computer where the Management service is installed appears on the Computers page.

The Configuration mode selection section appears.

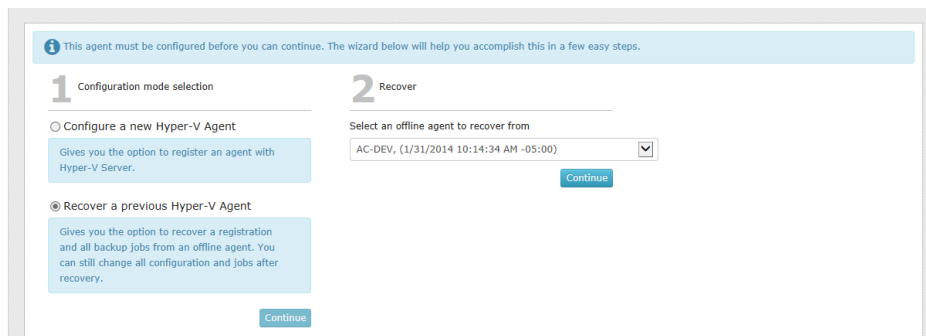
Note: The **Recover a previous Hyper-V Agent** option only appears if there is an offline Hyper-V Agent in the user's site. If the offline Hyper-V Agent was deleted from Portal but its data was not deleted from the vault, you can undelete the Hyper-V Agent. See [Undelete Hyper-V environments](#).



5. Select **Recover a previous Hyper-V Agent**, and then click **Continue**.

The **Recover** section appears. The **Select an offline agent to recover from** list shows the names of standalone Hyper-V host names and clusters where the Management service is offline, and shows the last date and time when the Management service connected to Portal.

Note: The date and time shown in the **Select an offline agent to recover from** list could reflect the date and time when the Management service was installed or the server was restarted. The date and time in this list does not reflect the date and time of the last backup.

The screenshot shows a configuration wizard with two main sections. The first section, '1 Configuration mode selection', has two radio button options: 'Configure a new Hyper-V Agent' and 'Recover a previous Hyper-V Agent'. The second section, '2 Recover', is active and contains a dropdown menu labeled 'Select an offline agent to recover from' with the selected value 'AC-DEV, (1/31/2014 10:14:34 AM -05:00)'. There are 'Continue' buttons at the bottom of both sections.

- From the **Select an offline agent to recover from** list, choose the protected environment from which you want to recover jobs and settings. If you are sure that this is the correct offline Agent, click **Continue**.

Note: Do not click **Continue** unless the correct offline Agent is selected. The offline Agent's settings and jobs are downloaded immediately after you click **Continue**.

The system downloads the offline agent's jobs and settings. If the offline agent was protecting a Hyper-V cluster, the name of the original protected cluster now appears on the Computers page instead of the Management service computer name. You cannot change the name to the name of the current cluster.

The **Success** section lists the passwords that you need to enter: Hyper-V environment, vault registrations, job encryption, application-consistent backups, and Email notifications.

- Click **Continue**.
- On the **Cluster Credentials** tab, do one of the following:
 - To continue protecting the same Hyper-V environment, enter the password for the specified user.
 - To provide credentials for a new Hyper-V environment so you can restore VMs to the new environment, enter Hyper-V environment information in the **Address** and **Domain** boxes. In the **Username** box, type the domain administrator account that is used to authenticate with the Hyper-V cluster or standalone host. In the **Password** box, type the password for the specified user. For more information, see [Change credentials or the network address for accessing Hyper-V](#).

To determine whether the credentials are valid, click **Verify Information**. If the credentials are valid, click **Okay** in the confirmation message box.

9. Click **Save**. In the confirmation message box, click **Okay**.
10. On the **Vault Settings** tab, enter the password for each vault connection. See [Add vault settings](#).
11. On the **Jobs** tab, edit each job and do the following:
 - In the **Encryption Password** and **Confirm Password** boxes, enter the job's data encryption password.
 - If application-consistent backups are enabled in the job, do one or both of the following:
 - To enter credentials for VMs in the job, enter an admin user's username and password in the Guest VM Credentials area.

The specified user must have admin access to VMs in the backup job. You can enter the username as `username` or `domain\username`.
 - To enter credentials for a specific VM in the job, click the arrow at the right side of the VM name in the **Protected by this job** box, and enter an admin user's username and password.

The specified user must have admin access to the VM. You can enter the username as `username` or `domain\username`.

If you enter credentials for a specific VM in the job, the Agent will not attempt to connect to the VM using the Guest VM Credentials.

IMPORTANT: If you do not enter credentials for VMs in the backup job, backups will be crash-consistent. Credentials are required for all application-consistent backups in Hyper-V environments— with or without log truncation.
- Click **Save**. In the confirmation message box, click **Continue**. See [Edit a Hyper-V backup job](#).
12. If required, on the **Advanced** tab, on the **Notifications** tab, enter the SMTP password. See [Set up email notifications for backups on a computer](#).
13. Click **Save**. In the confirmation message box, click **Okay**.
14. Reinstall the Host service on each Hyper-V cluster host or standalone node. See [Install the Hyper-V Agent Host service](#).

Note: If you reinstall the Management service in a Hyper-V environment for any reason, you must also reinstall each Host service.
15. If the protected Hyper-V VMs exist in the environment (i.e., the VMs were restored after a disaster or remained intact when the Hyper-V Agent was lost), you can re-enable all scheduled backup jobs for the Hyper-V environment. See [Disable or enable all scheduled backup jobs](#).

7.1 Hyper-V disaster recovery and migration

During a disaster recovery or when moving to a new Hyper-V environment, you can recover jobs and settings from an offline Hyper-V Agent. As described in [Recover jobs and settings from an offline Hyper-V Agent](#):

- the newly-installed Hyper-V Agent must be the same version or a later version than the offline agent.
- the new Hyper-V environment must be the same version or a later version than the environment protected by the offline agent.

Note: The following table outlines the process of recovering a protected Hyper-V environment when you have to replace one or more of the following components:

- Hyper-V Agent Management service

Note: You do not need to recover settings separately for a Host service that is lost or becomes unavailable. Host services upload their settings and logs to the Management service. When you register a Host service to a Management service, the Host service obtains its settings from the Management service.

- Hyper-V cluster or standalone host. This process can also be used when moving to a new Hyper-V environment.
- Portal

IMPORTANT: Configuration data, vault, and job information for the Hyper-V Agent is saved in the Portal database. To ensure that the Portal and a Hyper-V environment can be fully restored if the Portal is lost, the Portal database must be backed up. For more information, see the *Portal Installation and Configuration Guide*.

Component Lost			Recovery Process
Hyper-V Agent Management service	Hyper-V environment (cluster or standalone)	Portal	
✓			<ol style="list-style-type: none"> 1. Install the Hyper-V Agent Management service, and recover configuration information and jobs from the offline Hyper-V Agent. See Recover jobs and settings from an offline Hyper-V Agent. 2. If the IP address of the Hyper-V Agent Management server has changed, and Hyper-V Agent Host services were registered to the Management service using the IP address, do the following: <ol style="list-style-type: none"> a. Back up the host service log on each host where the Hyper-V Agent Host service is installed. The log is named AgentWorker.XLOG and is saved in a \Data\Logs subfolder in the Host service installation folder. <i>Note:</i> This step is suggested as a precaution, in case the host service log was not uploaded to the Management server. b. Uninstall and then reinstall each Hyper-V Agent Host service, and register each Host service to the Management service.
✓	✓		<ol style="list-style-type: none"> 1. Create a new Hyper-V cluster or standalone host (i.e., Windows server with Hyper-V role). 2. Install the Hyper-V Agent Management service in the new Hyper-V environment, and recover configuration information and jobs from the offline Hyper-V Agent. Enter information for the new Hyper-V environment (rebuilt in Step 1) on the Cluster Credentials tab. See Recover jobs and settings from an offline Hyper-V Agent. 3. Install the Hyper-V Agent Host service on each host, and register it to the Management service. See Install the Hyper-V Agent Host service. 4. Restore VMs. See Restore Hyper-V VMs.

Component Lost			Recovery Process
Hyper-V Agent Management service	Hyper-V environment (cluster or standalone)	Portal	
✓	✓	✓	<ol style="list-style-type: none"> 1. Restore the Portal and its protected database. See the <i>Portal Installation and Administration Guide</i>. 2. Rebuild the lost Hyper-V cluster or standalone host (i.e., Windows server with Hyper-V role). 3. Install the Hyper-V Agent Management service, and recover configuration information and jobs from the offline Hyper-V Agent. Enter information for the new Hyper-V environment (rebuilt in Step 1) on the Cluster Credentials tab. See Recover jobs and settings from an offline Hyper-V Agent. 4. Install the Hyper-V Agent Host service on each host, and register it to the Management service. See Install the Hyper-V Agent Host service. 5. Restore VMs. See Restore Hyper-V VMs.

8 Delete jobs and computers, and delete data from vaults

Regular users and Admin users can delete backup jobs from Portal without deleting associated data from vaults. See [Delete a backup job without deleting data from vaults](#). Admin users can delete computers and protected environments from Portal without deleting associated data from vaults. See [Delete a computer without deleting data from vaults](#).

In a Portal instance where the data deletion feature is enabled, Admin users can also:

- Delete backup jobs from Portal and submit requests to delete the job data from vaults. See [Delete a backup job and delete job data from vaults](#).

When deleting job data from vaults, there is a 72-hour waiting period before the data deletion request is sent to vaults. This waiting period gives Admin users in the site an opportunity to cancel the data deletion, if required. See [Cancel a scheduled job data deletion](#). During the waiting period, the job continues to run as scheduled.

- Delete computers from Portal and submit requests to delete the computer data from vaults. See [Delete a computer and delete computer data from vaults](#).

Note: Beginning in Portal 8.90, Admin users can submit requests to delete data from vaults for online or offline computers. In previous Portal versions, requests to delete data from vaults could only be submitted for online computers.

When deleting computer data from vaults, there is a 72-hour waiting period before the data deletion request is sent to vaults. This waiting period gives Admin users in the site an opportunity to cancel the data deletion, if required. See [Cancel a scheduled computer data deletion](#). During the waiting period, the computer's jobs continue to run as scheduled.

8.1 Delete a backup job without deleting data from vaults

Regular users and admin users can delete backup jobs from online computers without deleting the job data from vaults. Because the data remains in the vaults, you will be billed for it.

For a Hyper-V Agent, if a job is deleted from Portal but the job data is not deleted from vaults, the data can be restored using the regular restore procedure. See [Restore Hyper-V VMs](#).

In a Portal instance where the data deletion feature is enabled, Admin users can submit requests to delete job data from vaults when they delete jobs from Portal. See [Delete a backup job and delete job data from vaults](#).

To delete a backup job without deleting data from vaults:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the online computer with the job that you want to delete, and expand its view by clicking its row.
3. Click the **Jobs** tab.

4. In the **Select Action** menu of the job that you want to delete, click **Delete Job**.
5. If you are signed in as an Admin user in a Portal instance where the data deletion feature is enabled, a Delete Job dialog box appears.

To delete the backup job without deleting data from vaults, click **Remove job** and then click **Delete**.

Note: The Delete Job dialog box does not appear if you cannot delete backup data in vaults because your Portal instance does not support vault data deletion or you are signed in as a regular user.

6. In the confirmation dialog box, type **CONFIRM**.

Note: You must type **CONFIRM** in capital letters.

7. Click **Confirm Deletion**.

8.2 Delete a backup job and delete job data from vaults

In a Portal instance where the data deletion feature is enabled, Admin users can delete backup jobs and request that data for the jobs be deleted from all vaults. After the data is deleted from the vaults, you will not be billed for it.

To protect against inadvertently deleting the wrong data, the data deletion is scheduled for 72 hours after the request is made and an email notification is sent to Admin users in the site and to Super users. During the waiting period, the job continues to run as scheduled.

During the 72-hour waiting period before job data is deleted, Admin users can cancel scheduled job data deletions in their sites. See [Cancel a scheduled job data deletion](#).

If a scheduled job data deletion is not canceled during the 72-hour waiting period, the job is deleted from Portal, the deletion request is sent to vaults and job data is automatically deleted from associated vaults. If data for a job cannot be deleted for some reason, an email notification is sent to a vault administrator. The vault administrator can then manually delete the data.

Note: Because the data is available for restore during the 72-hour waiting period, it will still be included in customer invoices. Usage reduction for invoicing purposes does not occur until the data is deleted.

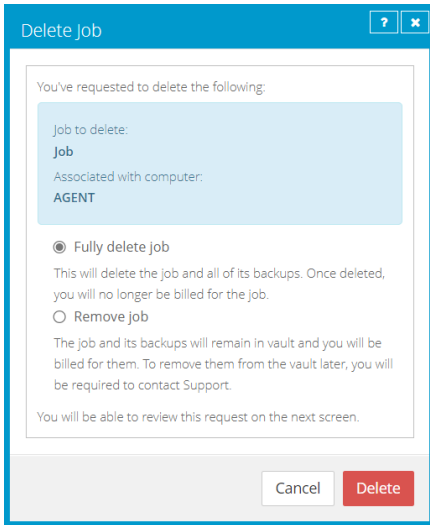
WARNING: Job data deletion is permanent. After the data is deleted from vaults, it cannot be recovered or restored.

To delete a backup job and delete job data from vaults:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Find the computer with the job that you want to delete, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the **Select Action** menu of the job that you want to delete, click **Delete Job**.

A Delete Job dialog box appears if the data deletion feature is enabled in your Portal instance.

Note: If the Delete Job dialog box does not appear, you cannot request that data for the job be deleted from vaults. You can only delete the job from Portal. See [Delete a backup job without deleting data from vaults](#).



5. Select **Fully delete job**, and then click **Delete**.

IMPORTANT: To permanently delete unnecessary data from vaults and reduce billing, you must select **Fully delete job**. If you select **Remove job**, data will not be removed from vaults and your invoice will not be affected.

WARNING: Job data deletion is permanent. After the data is deleted from vaults, it cannot be recovered or restored.

6. In the confirmation dialog box, type **CONFIRM**.

Note: You must type **CONFIRM** in capital letters.

7. Click **Confirm Deletion**.

A Job Deleted dialog box states that the job and associated data in your vaults is scheduled to be deleted.

8. Click **Close**.

The Last Backup Status column shows **Scheduled For Deletion** for the job. The Date column shows the date when the job will be deleted from Portal and job data will be deleted from vaults. Within a day of the scheduled deletion, the Date column will also show the time when the job and its data will be deleted.



Beginning in Portal 9.10, when a job is scheduled for deletion, the **Scheduled for Deletion** status appears for every instance of the job in Portal. A job can appear for multiple computers if a computer was re-registered or the Restore from Another Computer workflow was used.

An email is sent to Admin users in the site and to Super users to indicate that the job deletion has been scheduled. During the 72-hour waiting period before data is deleted, you can cancel the deletion request. Because the data is available for restore during this period, it will still be included in customer invoices. Usage reduction for invoicing purposes does not occur until the data is deleted.

When data deletion is in progress for a job, the **Deletion in Progress** status appears for the job. Beginning in Portal 9.20, the **Scheduled for Deletion** status appears for every instance of the job in Portal.

When a job is deleted from vaults, the job is deleted from all computers where it appears.

8.3 Cancel a scheduled job data deletion

In a Portal instance where the data deletion feature is enabled, Admin users can delete a backup job and request that data for the job be deleted from all vaults. The data deletion is scheduled for 72 hours after the request is made and an email notification is sent to Admin users in the site and to Super users.

During the 72-hour period before a job is deleted from Portal and the job data is deleted from vaults, Admin users in the site can cancel the data deletion. If a scheduled data deletion is canceled, an email notification is sent to Admin users in the site and to Super users.


Beginning in Portal 9.10, when a job is scheduled for deletion, the **Scheduled for Deletion** status appears for every instance of the job in Portal. A job can appear for multiple computers if a computer was re-registered or the Restore from Another Computer workflow was used. An Admin user can cancel the deletion from any instance of the job.

To cancel a scheduled job data deletion:

1. When signed in as an Admin user, click **Computers** on the navigation bar.

The Computers page shows registered computers.

2. Find the computer with the scheduled job data deletion that you want to cancel, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the Select Action menu of the job that is scheduled for deletion, click **Cancel Delete**.



A confirmation dialog box asks whether you want to cancel the deletion.

5. Click **Yes**.

Values in the Last Backup Status and Date columns for the job revert to the values that appeared before the job was scheduled for deletion.

An email is sent to Admin users in the site and to Super users to indicate that the scheduled job deletion has been canceled.



8.4 Delete a computer without deleting data from vaults

Admin users can delete computers from Portal without deleting the computer data from vaults. You can delete both online and offline computers from Portal without deleting data from vaults. Because the data remains in the vaults, you will be billed for it.

If you delete an offline Hyper-V environment from Portal, you must undelete the environment before you can recover jobs and settings from the Hyper-V Agent. See [Undelete Hyper-V environments](#) and [Recover jobs and settings from an offline Hyper-V Agent](#).

Note: When a computer is deleted from Portal, the agent is not removed from the computer where it is installed. To remove the agent from the computer, you must manually uninstall it.

To delete a computer without deleting data from vaults:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Select the check box for each computer that you want to delete.
3. In the **Actions** list, click **Delete Selected Computer(s)**.

4. If the data deletion feature is enabled in your Portal instance, a Delete Computer(s) dialog box appears.

To delete the computer without deleting data from vaults, click **Remove computer(s) from Portal only** and then click **Delete**.

Note: The Delete Computer(s) dialog box only appears if your Portal instance supports vault data deletion.

5. In the confirmation dialog box, type **CONFIRM**.

Note: You must type **CONFIRM** in capital letters.

6. Click **Confirm Deletion**.
7. In the confirmation dialog box, click **Yes**.
8. In the Success dialog box, click **Okay**.


8.5 Undelete Hyper-V environments

By default, the Computers page in Portal shows computers that are registered to Portal and are online, offline, or online and need to be rebooted.

You can choose a different view on the Computers page to see protected Hyper-V environments that have been deleted from Portal. You can also “undelete” deleted Hyper-V environments so you can recover data from the environments. See [Recover jobs and settings from an offline Hyper-V Agent](#).

To undelete a Hyper-V environment:

1. On the navigation bar, click **Computers**. The Computers page shows registered computers.
2. Click the views list at the top of the page.



3. In the views list, click the **Deleted** view.

The Computers page shows Hyper-V environments that have been deleted from Portal.

4. Select the check box for each Hyper-V environment that you want to undelete.
5. In the confirmation dialog box, click **Yes**.
6. In the Success dialog box, click **Okay**.

8.6 Delete a computer and delete computer data from vaults

In a Portal instance where the data deletion feature is enabled, Admin users can delete computers and request that data for the computers be deleted from all vaults. After the data is deleted from the vaults, you will not be billed for it.

Note: Beginning in Portal 8.90, Admin users can submit requests to delete data from vaults for online or offline computers. In previous Portal versions, requests to delete data from vaults could only be submitted for online computers.

To protect against inadvertently deleting the wrong data, the data deletion is scheduled for 72 hours after the request is made, an email notification is sent to Admin users in the site and to Super users, and the status of the computer in Portal changes to *Scheduled for deletion*. During the waiting period, the computer's jobs continue to run as scheduled.

During the 72-hour waiting period before a computer data deletion request is sent to vaults, Admin users in the site can cancel the scheduled computer data deletion. See [Cancel a scheduled computer data deletion](#).

If a scheduled computer data deletion is not canceled during the 72-hour waiting period, the deletion request is sent to vaults and job data is automatically deleted from associated vaults. If data for a computer

cannot be deleted for some reason, an email notification is sent to a vault administrator. The vault administrator can then manually delete the data. After the computer data is deleted from vaults, the computer is deleted from Portal.

Note: Because the data is available for restore during the 72-hour waiting period, it will still be included in customer invoices. Usage reduction for invoicing purposes does not occur until the data is deleted.

Note: When a computer is deleted from Portal, the agent is not removed from the computer where it is installed. To remove the agent from the computer, you must manually uninstall it.

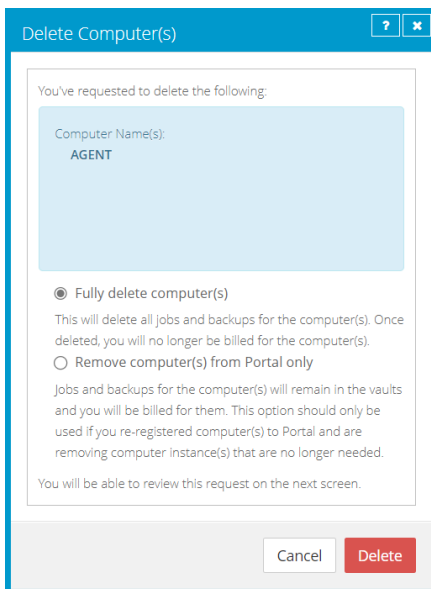
WARNING: Computer data deletion is permanent. After the data is deleted from vaults, it cannot be recovered or restored.

To delete a computer and delete computer data from vaults:

1. When signed in as an Admin user, click **Computers** on the navigation bar.
The Computers page shows registered computers.
2. Select the check box for each computer that you want to delete.
3. In the **Actions** list, click **Delete Selected Computer(s)**.

A Delete Computer(s) dialog box appears if the data deletion feature is enabled in your Portal instance.

Note: If the Delete Computer(s) dialog box does not appear or the **Fully delete computer(s)** option is not available, you cannot request that data for the selected computers be deleted from vaults. You can only delete the selected computers from Portal. See [Delete a computer without deleting data from vaults](#).



4. Select **Fully delete computer(s)**, and then click **Delete**.

IMPORTANT: To permanently delete unnecessary data from vaults and reduce billing, you must select **Fully delete computer(s)**. If you select **Remove computer(s) from Portal only**, data will not be removed from vaults and your invoice will not be affected.

WARNING: Computer data deletion is permanent. After the data is deleted from vaults, it cannot be recovered or restored.

5. In the confirmation dialog box, type **CONFIRM**.

Note: You must type **CONFIRM** in capital letters.

6. Click **Confirm Deletion**.

WARNING: Computer data deletion is permanent. After the data is deleted from vaults, it cannot be recovered or restored.

A Computer(s) Deleted dialog box states that the computer(s) and associated data in your vault(s) are scheduled to be deleted.

7. Click **Close**.

The Status column shows *Scheduled for deletion* for the computer(s). If you expand the computer, a message indicates when the computer is scheduled to be deleted.

During the 72-hour period, you can cancel the deletion request. Because the data is available for restore during this period, it will still be included in customer invoices. Usage reduction for invoicing purposes does not occur until the data is deleted.

You cannot add, edit, run, schedule or delete jobs for a computer that is scheduled for deletion. Existing backup jobs continue to run as scheduled until the computer is deleted.

8.7 Cancel a scheduled computer data deletion

In a Portal instance where the data deletion feature is enabled, Admin users can delete an online computer and request that data for the computer be deleted from all vaults. The data deletion is scheduled for 72 hours after the request is made. See [Delete a computer and delete computer data from vaults](#).

During the 72-hour period before a computer data deletion request is set to vaults, Admin users in the site can cancel the data deletion. If a scheduled data deletion is canceled, an email notification is sent to Admin users in the site and to Super users.

To cancel a scheduled computer data deletion:

1. When signed in as an Admin user, click **Computers** on the navigation bar.

The Computers page shows registered computers.

2. Select the check box for each computer for which you want to cancel the scheduled data deletion.

The Status column shows *Scheduled for deletion* for each computer that is scheduled for deletion.

3. In the Actions list, click **Cancel Deletion of Selected Computers**.

Note: If **Cancel Deletion of Select Computers** is not available, the data deletion request for a selected computer may have already been sent to vaults. To see when a computer was scheduled for deletion, expand the computer row.

A confirmation dialog box asks whether you want to cancel the deletion.

4. Click **Yes**.

A Success dialog box appears.

5. Click **Okay**.

The value in the Status column for each computer reverts to the value that appeared before the computer was scheduled for deletion.

An email is sent to Admin users in the site and to Super users to indicate that the scheduled computer deletion has been canceled.

9 Monitor computers, jobs and processes

You can monitor backups, restores and protected computers using the following features in Portal:

- **Current Snapshot.** The Current Snapshot provides total numbers of backups and computers in various categories in your site, and allows you to navigate to more detailed information. See [Monitor backups and computers using the Current Snapshot](#).
- **Site Usage charts.** In Portal instances that obtain data from billing systems, a Site Usage chart can show the amount of data backed up for a site in a billing period compared to a usage checkpoint amount. See [Monitor storage usage using Site Usage charts and emailed alerts](#).
- **Computers page.** The Computers page shows status information for computers and their jobs. See [View computer and job status information](#). You can also access logs for unconfigured computers from this page. See [View an unconfigured computer's logs](#).
- **Process Details dialog box.** This dialog box shows information about all running, queued and recently-completed processes for a job. See [View current process information for a job](#).
- **Email notifications.** To make it easier to monitor backups, users can receive emails when backups finish or fail. See [Set up email notifications for backups on a computer](#).
- **Process logs and safeset information.** Process logs indicate whether each backup and restore completed successfully, and provide information about any problems that occurred. You can also view information about the safeset created by a specific backup. See [View a job's process logs and safeset information](#) and [View a Hyper-V VM's backup history and logs](#).
- **Monitor page.** The Monitor page shows the most recent backup status for each job, and allows you to navigate to the computer and job for each backup. See [View, export and email backup statuses on the Monitor page](#).

9.1 Monitor backups and computers using the Current Snapshot

In the Current Snapshot on the Dashboard, you can view total numbers of backup jobs and computers in your site in various categories. You can then navigate from these totals to view more detailed information about the jobs and computers.

To monitor backups and computers using the Current Snapshot:

1. On the navigation bar, click **Dashboard**.

The Current Snapshot at the left side of the Dashboard shows the number of backup jobs and computers in the following categories:

- **Backups Requiring Attention** — Number of backup jobs where the last backup attempt failed, completed with errors, did not back up any files, reached a license limit, was cancelled or had a potential ransomware threat.
- **Missed Backups** — Number of backup jobs that have not run for seven days.

- **Backups With Warnings** — Number of backup jobs where the last backup attempt completed with warnings, was deferred, was deferred with warnings or was skipped. This category also includes backup jobs that have never run.
 - **Computers Requiring Reboot** — Number of computers with a pending reboot.
 - **Offline Computers** — Number of computers that are not currently in contact with Portal. A computer can be offline if it is turned off, if the Agent has been uninstalled from the system, or if the system no longer exists.
 - **Computers Scheduled for Deletion** — Number of computers that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
 - **Computers With Certificate Failures** — Number of computers reporting a certificate failure. See [Resolve certificate failures](#).
 - **Total Computers** — Total number of computers in the site.
 - **Successful Backups** — Number of backup jobs where the last backup attempt completed without errors, warnings, or deferrals.
 - **Jobs Scheduled for Deletion** — Number of jobs that are scheduled for deletion from Portal and from vaults. This category is only applicable to Portal instances where the data deletion feature is enabled.
2. To view computers in a particular site, click the sites box in the top right of the Current Snapshot box. In the menu, click the site that you want to view.
Computers in the selected site appear on the Computers page.
 3. To view information about backup jobs or computers in one of the categories, click the category.
If you click **Potential Threats**, **Backups Requiring Attention**, **Missed Backups**, **Backups With Warnings** or **Successful Backups**, backup jobs in the category appear on the Monitor page.
If you click **Computers Requiring Reboot**, **Offline Computers**, **Computers Scheduled For Deletion**, **Computers With Certificate Failures** or **Total Computers**, computers in the category appear on the Computers page.

9.2 Monitor storage usage using Site Usage charts and emailed alerts

For some sites in some Portal instances, Admin users can view a Site Usage chart on the Dashboard. This chart shows the amount of data backed up for computers in the site compared to a specified limit. This can help customers monitor their storage usage and avoid billing overages.

When this feature is enabled for a site, Admin users for the site also receive email alerts when the site's storage usage first reaches 50%, 75%, 90% and 100% of the specified limit. If a site's storage usage is above 50%, 75%, 90% or 100% of the specified limit at the start of a billing period, Admin users also receive an email alert at the start of the billing period. Admin users cannot opt out of usage email alerts when this feature is enabled.

Note: At the start of a billing period, Portal might show usage data and send email alerts for the previous billing period. Usage data and alerts are provided for the new billing period as soon as the data is available.

Site Usage charts and emailed alerts are available beginning in Portal 9.30 in some Portal instances that obtain data from billing systems. Admin users in a Parent site can enable this feature and specify a limit or "User Checkpoint" for eligible sites.


Support users can also view Site Usage charts.

To monitor storage usage using Site Usage charts:

1. Sign in to Portal as an Admin user.
2. On the navigation bar, click **Dashboard**.

If usage tracking and alerting is enabled for your site, a Site Usage chart appears at the right side of the Dashboard. The chart shows the amount of data backed up for computers in the site in the current billing period as compared to the specified limit, or usage checkpoint amount. The amount of data backed up is the original size of the data before it was compressed.

If you are viewing a parent site, a separate Site Usage chart could appear for the parent site and any child sites where this feature is enabled. If more than four charts appear, you can scroll through the charts.



As described in the table below, the Site Usage chart color indicates how much data has been backed up in the current billing period compared to the specified limit, or usage checkpoint amount:

Chart color	Description
Green	The site's storage usage in the current billing period is less than 50% of the specified limit.
Yellow	The site's storage usage in the current billing period is between 50% and 75% of the specified limit.
Orange	The site's storage usage in the current billing period is between 75% and 90% of the specified limit. An orange warning message appears beside the chart in this case.
Red	The site's storage usage in the current billing period is more than 90% of the specified limit. A red warning message appears beside the chart in this case.

If a Site Usage chart does not appear, usage tracking and alerting might not be available for your site or in your Portal instance.

9.3 View computer and job status information

On the Computers page in Portal, you can view status information for computers and their jobs.







To view computer and job status information:


1. On the navigation bar, click **Computers**.



The Computers page shows registered computers.

The Availability column indicates whether each computer is online or offline. Online computers are in contact with Portal, while offline computers are not currently available. A computer can be offline if it is turned off, if the agent has been uninstalled from the system, or if the system has been lost.

The Status column shows the status of each computer. Possible statuses include:




-  OK — Indicates that all jobs on the computer ran without errors or warnings.
 -  OK with warnings — Indicates that one or more of the computer’s jobs completed with warnings.
 -  Attention — Indicates that one or more of the computer’s jobs failed or completed with errors.
 -  Unconfigured — Indicates that no jobs have been created for the computer.
 -  Scheduled for deletion — Indicates that the computer is scheduled for deletion from Portal and from vaults. This status only appears in Portal instances where the data deletion feature is enabled.
 -  Certificate failure — Indicates that the agent is reporting a certificate change.
2. Find the computer for which you want to view status information, and click the row to expand its view.
 3. View the **Jobs** tab.

If a backup or restore is running for a job, a Process Details symbol  appears beside the job name, along with the number of processes that are running.

Name	Job Type
 1 job1	Local System
 1 job2	Local System


If you click the Process Details symbol, the Process Details dialog box shows information about processes for the job. See [View current process information for a job](#).

The **Last Backup Status** column shows the last backup status reported for each job. An agent reports a backup status to Portal each time it starts, skips or completes a backup. Possible statuses include:








-  Completed — Indicates that the last backup completed successfully, and a safeset was created.
-  Completed with warnings — Indicates that the last backup completed and a safeset was created, but problems occurred during the backup. For example, a warning could indicate that a file or volume that was selected in the backup job was not available for backup.
-  Deferred — Indicates that the last backup was deferred. A safeset was created, but not all data that was selected was backed up.

Deferring is used to prevent large backups from running at peak network times. When deferring is enabled, a backup job does not back up any new data after a specified amount of time.

Hyper-V VM backups can be deferred when they are run manually (ad hoc), but not when they are scheduled to run.

-  Skipped — Indicates that a backup was skipped. Backups are sometimes skipped if they are scheduled to run multiple times per day.

Hyper-V Agent backups cannot be skipped.

-  Never Run — Indicates that the backup job has never run.
-  Missed — Indicates that the job has not run for 7 days.
-  Completed with errors — Indicates that the backup completed and a safeset is available for restore, but problems occurred. Typically, this status indicates that not all of the data was backed up. For Hyper-V environments, this status can appear when problems were encountered during the backup, but the backups later succeeded.
-  No Files backed up — Indicates that no files were backed up during the last backup attempt
-  Failed — Indicates that the backup failed and no safeset was created. For Hyper-V backups, this status indicates that the backup failed for all virtual machines in the job.
-  Cancelled — Indicates that the backup was cancelled.
-  Scheduled for Deletion — Indicates that the job is scheduled to be deleted from Portal and job data is scheduled to be deleted from all vaults on the date shown in the Date column. This

backup status is only possible in Portal instances where the data deletion feature is enabled. See [Delete a backup job and delete job data from vaults](#).

To view logs for a job, click the job status. For more information, see [View a job’s process logs and safeset information](#).

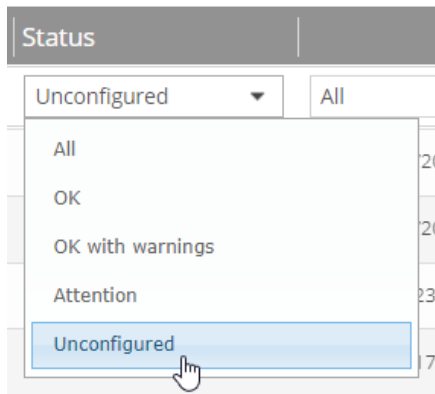
9.4 View an unconfigured computer’s logs

You can view logs for unconfigured computers that are online. Unconfigured computers do not have any backup jobs.

To view an unconfigured computer’s logs:

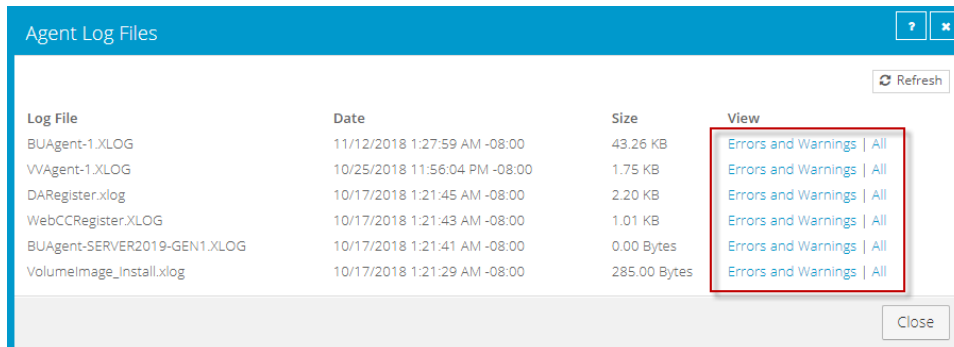
1. On the navigation bar, click **Computers**.

The Computers page shows registered computers. To only show unconfigured computers, click “Unconfigured” in the **Status** filter.



2. Find an unconfigured computer that is online, and expand its view by clicking the computer row.
3. Click the **logs** link for the unconfigured computer.

The Agent Log Files window shows a list of logs for the computers. Links to the logs appear at the right side of the window.



4. Do one of the following:


- To only view errors and warnings in a log, click **Errors and Warnings** for the log.
- To view an entire log, click **All** for the log.

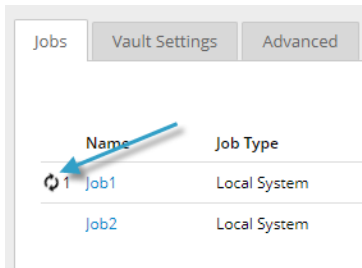
The log appears in a new browser tab.


9.5 View current process information for a job


In the Process Details dialog box, you can view information about running, queued and recently-completed processes for a job. Processes include backups, restores, and synchronizations, and is typically deleted within an hour after the process ends.

To view current process information for a job:

1. While a backup, restore, or synchronization is running, do one of the following:
 - On the Computers page, on the Jobs tab, click the Process Details symbol  beside the job name.

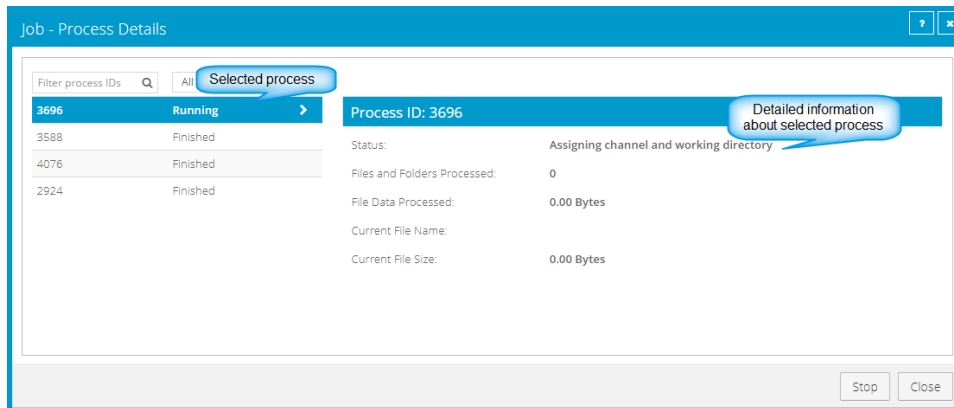


- On the Monitor page, click the Process Details symbol  beside the job name.



Note: For a Hyper-V restore, a Process Details symbol does not appear on the Monitor page. Instead, a line for the Hyper-V environment appears with "Restoring Virtual Machine" as the Last Backup Status. To view the environment's running processes, click the Hyper-V environment name, and then click its Virtual Machines tab on the Computers page.


If you clicked a Process Details symbol, the Process Details dialog box lists backup, restore and synchronization processes that are running, queued and recently completed for the job. Detailed information is shown for the process that is selected on the left side of the dialog box.



- To view information about a different process, click the process or VM name on the left side of the dialog box.

Detailed information is shown at the right side of the dialog box.

- If the Process Details dialog box lists backup, restore and synchronization processes for the job, do one of the following in the status list to show only some processes:
 - To only show queued processes, click **Launched**.
 - To only show processes that are waiting for user action, click **Operator Request**.
 - To only show processes that are in progress, click **Running**.
 - To only show completed processes, click **Finished**.
 - To only show processes that are finishing, click **Finalizing**.



9.6 View a job’s process logs and safeset information

To determine whether a backup, restore or other process completed successfully, or to determine why a process failed, you can view a job’s process logs.

You can also view information about safesets created for the job. A safeset is an instance of backup data on the vault.

To view a job’s process logs and safeset information:


- On the navigation bar, click **Computers**.
The Computers page shows registered computers.
- Find the computer for which you want to view logs, and click the row to expand its view.

On the **Jobs** tab, the **Last Backup Status** column shows the status of each backup job.

3. To view log files for a job, do one of the following:

- In the job's **Select Action** menu, click **Logs**.
- In the **Last Backup Status** column, click the job status.

The Logs window lists the most recent backups, restores and other processes in the Hyper-V environment.

4. To view processes for a different day, click the calendar button.  In the calendar that appears, click the date of the log that you want to view.

5. In the list of processes on the selected date, click the process for which you want to view the log.

The window shows the selected log.

9.7 View, export and email backup statuses on the Monitor page

You can view recent job backup statuses on the Monitor page in Portal and navigate to related information on the Computers page or in the Logs window.

You can export data from the Monitor page in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format. The exported data file (named "Job Monitor Export.csv", "Job Monitor Export.xls" or "Job Monitor Export.pdf") is downloaded to the user's computer.

Beginning in Portal 9.20, Admin users and Support users can email reports with data from the Monitor page. These Job Monitor Export reports can be:

- Emailed once to one or more recipients. To specify which job backup statuses appear in this report, you can select a view and filter data on the Monitor page.
- Scheduled to be emailed to one or more recipients on specified days at a specified time. To specify which job backup statuses appear in a scheduled report, you can filter data by any column except the Last Backup Date column. You can only schedule a report to be emailed from the All Jobs view on the Monitor page.

A Job Monitor Export report is emailed as an attachment in .csv, .xls or .pdf format (named "Job Monitor Export.csv", "Job Monitor Export.xls" or "Job Monitor Export.pdf"). Reports in .xls and .pdf format are formatted using the site's logo, color, and custom text.

Note: We recommend turning off macros in Microsoft Excel when using Portal, particularly if you export or email information in .xls or .csv format and open these reports in Excel.


To view, export and email backup statuses on the Monitor page:

1. On the navigation bar, click **Monitor**.

The Monitor page shows recent backup statuses for jobs in your site.

2. To change which job backup statuses appear, click a view or enter filter criteria.

3. To view information for a job or computer on the Computers page, click the name of a job or online computer.
4. To view a job's logs in the History/Logs window, click the job's last backup status.
5. To export job backup status data from the Monitor page, click the **Export** box. In the list that appears, click one of the following formats for the exported data file:
 - CSV (comma-separated values)
 - XLS (Microsoft Excel)
 - PDF (Adobe Acrobat)



The data file is downloaded to your computer in the specified format.


6. To email a Job Monitor Export report, do the following when signed in as an Admin or Support user:
 - a. To specify which job backup statuses appear in the report, click a view or enter filter criteria.
 - b. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Email Once**.
 - c. In the Email Once dialog box, do the following:
 - i. In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - ii. In the **Subject** box, type a subject for the report email.
 - iii. In the Attachment list, click one of the following formats for the emailed report:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)
 - d. Click **Okay**.
7. To schedule a Job Monitor Export report to be emailed, do the following when signed in as an Admin or Support user:
 - a. To specify which job backup statuses appear in the scheduled report, enter filter criteria in any column except the Last Backup Date column.

Note: You can only schedule a report to be emailed when the All Jobs view is selected on the Monitor page.

- b. Click the **Email/Schedule** box. In the **Send Report** list that appears, click **Schedule New Report**.
- c. In the Email/Schedule dialog box, do the following:
 - In the **To** box, type one or more email addresses for sending the report. Use commas to separate multiple email addresses.
 - In the **Report Name** box, type a name for the scheduled report. This name appears in the **Email/Schedule** list.
 - In the **Subject** box, type a subject for the email.
 - In the **Attachment** list, click one of the following formats for the emailed report:
 - CSV (comma-separated values)
 - Excel (Microsoft Excel)
 - PDF (Adobe Acrobat)
- d. Do one of the following:
 - To email the report on specific days each week, in the **Frequency** list, click **Daily**. In the day row, select the days when you want to email the report each week.




- To email the report once each week, in the **Frequency** list, click **Weekly**. In the day row, select the day when you want to email the report each week.



- To email the report once each month, in the **Frequency** list, click **Monthly**. In the calendar, select the date when you want to email the report each month, or select

Last Day to email the report on the last day of each month.



- e. Using the **At** field, specify the time when you want to email the report on the specified days.
- f. Click **Okay**.








9.8 View a Hyper-V VM’s backup history and logs

Hyper-V backup jobs can include multiple VMs, but each VM is backed up as a separate task on the vault. You can view historical backup information and logs separately for each Hyper-V VM.

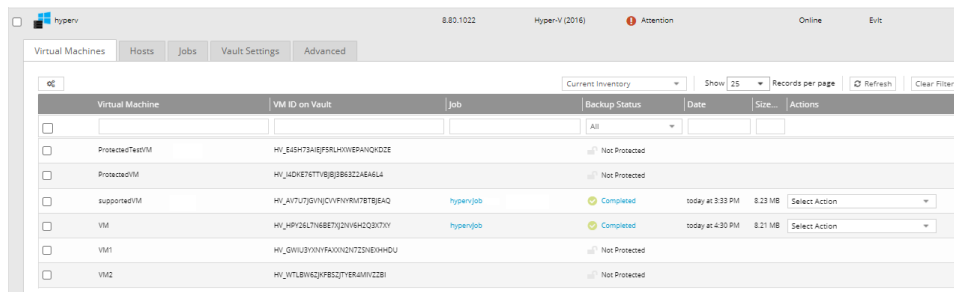
To view a Hyper-V VM’s backup history and logs:

1. On the navigation bar, click **Computers**.
The Computers page shows registered computers.
2. Find the Hyper-V environment for which you want to view the backup history and logs, and click the row to expand its view.
3. Click the **Virtual Machines** tab.

The **Virtual Machines** tab shows VMs in the Hyper-V cluster or standalone host. The **Backup Status** column shows the backup status of each VM. Possible statuses include:

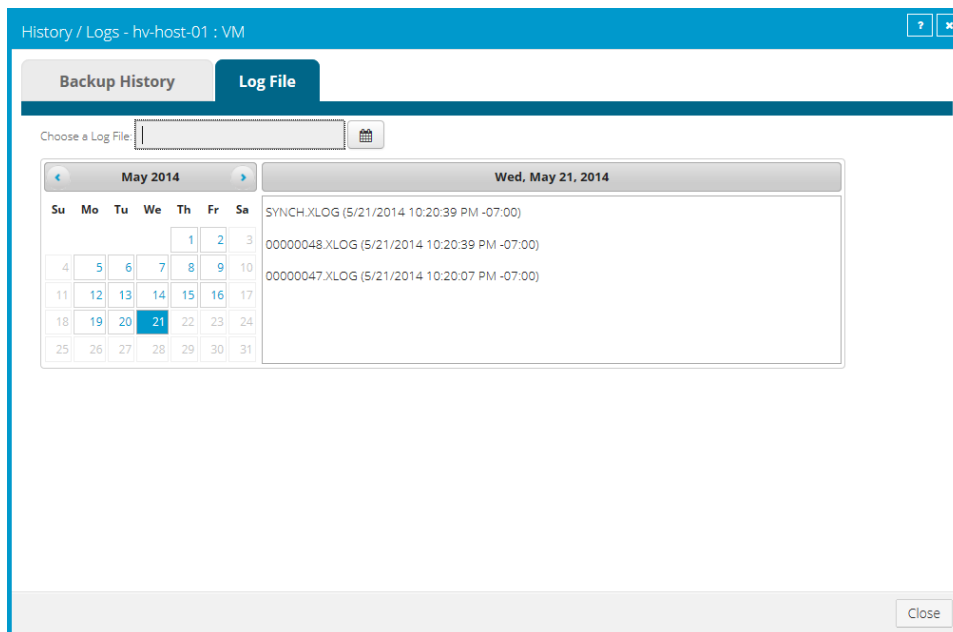
-  Completed — Indicates that the VM has been backed up.
-  Missed
-  Deferred
-  Not Protected — Indicates that the VM is not part of a backup job.
-  Never Run — Indicates that the VM is part of a backup job that was never run.
-  In Progress
-  Failed


-  Cancelled




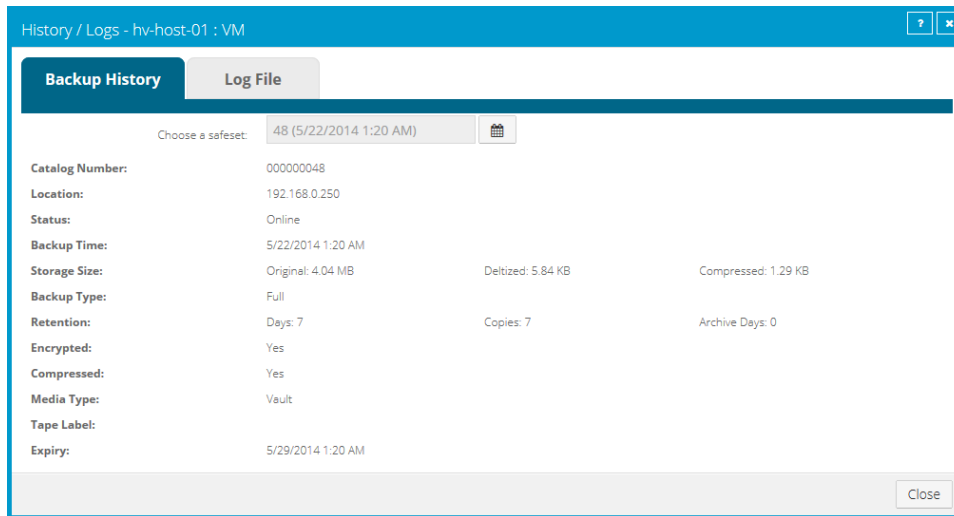
Virtual Machine	VM ID on Vault	Job	Backup Status	Date	Size	Actions
ProtectedTestVM	HY_E4SH73AIEJF8LKHVWEPANQDCE		Not Protected			
ProtectedVM	HY_H4DH676TVI8J2B932246464		Not Protected			
supportedVM	HY_A1V7JGVIVJCVFVYRUT8TBEAQ	hyperJob	Completed	today at 8:33 PM	8.23 MB	Select Action
VM	HY_HPVZL7N8BETJ2VWEHQ287XY	hyperJob	Completed	today at 4:30 PM	8.21 MB	Select Action
VM1	HY_GWU3YNYFA0NDZ73NEHMDU		Not Protected			
VM2	HY_WTLBIV62KFB5ZTYERAMVZBEI		Not Protected			

4. Click the **Backup Status** column of the VM for which you want to view the backup history and logs. The **History / Logs** window lists log files from the date selected in the calendar.



5. To view a log file for a process on the selected date, click the process. The log file appears.
6. To view a log file for a different date, click the calendar button.  In the calendar that appears, click the date of the log that you want to view. In the list of processes on the selected date, click the process for which you want to view the log. The log file appears.
7. To view safeset information for a particular VM backup, click the **Backup History** tab. The tab shows information for the VM's most recent backup.

To view information for a different safeset, click the calendar button.  In the calendar that appears, click the date of the backup for which you want to view information. In the list of backups on the selected date, click the backup for which you want to view information. The tab shows safeset information for the selected backup.



9.9 Hyper-V Agent logs and configuration files

Hyper-V Agent logs are saved in the Hyper-V Agent Management service installation folder, in a Data subfolder.

This folder includes logs from both the Hyper-V Agent Management and Host services. Host services upload logs to the Management service after a process ends.

Note: Because Hyper-V Agent Host services perform backups and restores, and do not upload logs until a process is completed, you cannot view backup or restore logs on the Management service computer while processes are running.

Because the `<ManagementServiceInstallFolder>\Data` folder also contains Hyper-V Agent configuration information, it can provide all information necessary for troubleshooting Agent issues. If information is required for troubleshooting, you can compress the `<ManagementServiceInstallFolder>\Data` folder as a .zip file and send it to your service provider.

10 Understanding and troubleshooting Hyper-V processes

This section provides information that can be helpful when monitoring and troubleshooting Hyper-V Agent processes.

10.1 Some VMs backed up before others

If a Hyper-V backup job includes more than one VM, each VM is backed up as a separate job or task on the vault. If you stop a Hyper-V backup job that is running, VMs in the job that have already been completely backed up remain on the vault.

In Hyper-V on Windows Server 2012 R2, the Hyper-V Agent first backs up VMs with virtual disks on a single cluster shared volume (CSV). Any VMs which span multiple CSVs are backed up in a separate stage after all VMs which have their virtual disks on a single CSV.

10.2 VM skipped during backup

A VM could be skipped during a backup for one of the following reasons:

- The VM's storage is being moved while the backup job is running. If you try to back up a VM during storage migration, the VM is skipped during the backup. Other VMs in the same job will be backed up.
- The VM shares a virtual hard disk. The Hyper-V Agent does not back up VMs that contain shared virtual hard disks. Shared virtual hard disks became available in Windows Server 2012 R2.
- In Hyper-V on Windows Server 2012 R2, the VM contains mixed storage (e.g., one virtual disk on local storage and another virtual disk on a CSV). Hyper-V Agent 9.10 does not back up VMs on Windows Server 2012 R2 that contain mixed storage. Other VMs in the same job are backed up.

10.3 VM backup fails

In Hyper-V on Windows Server 2012 R2, if a VM is stored on a CSV that larger than 63 TB, backups for the VM could fail. Although Microsoft supports CSVs up to 64 TB in size, backups cannot finish successfully unless space is available for checkpoints (snapshots). For this reason, the supported volume size for the Hyper-V Agent is 64 TB minus 1024 GB (63 TB).

10.4 Live migration fails

If a VM is currently being backed up, live migration could fail for the VM.

10.5 VM restore fails

VMs that were backed up in a newer Hyper-V environment cannot be restored to an older Hyper-V environment. For example, you cannot restore VMs that were backed up in a Windows Server 2019 environment to a Windows Server 2016 environment).

11 Carbonite Server Backup Support

If you have a question about Carbonite Server Backup that isn't covered in this guide, our frequently-updated Knowledge Base contains comprehensive information. The Knowledge Base is your first stop when searching for any Carbonite Server Backup solutions you may need. We highly recommend searching here first for the quickest answers to your questions.

Knowledge Base: <http://support.carbonite.com/evault>

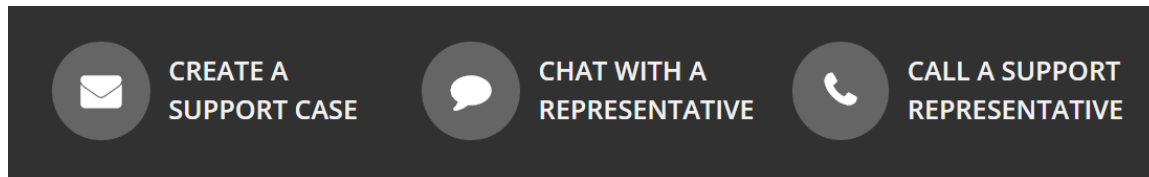
What can we help you with?

Popular Searches
[pending reboot](#), [restore](#), [clnt-e-04103](#)

11.1 Contacting Carbonite

If you need live assistance from a qualified support agent, Carbonite Support is here for you 24 hours a day, 7 days a week (excluding US holidays). Please feel free to get in touch with us, and we'll help out any way we can! You can find the contact information for Carbonite Support in the Knowledge Base:

<http://support.carbonite.com/evault>



A dark grey horizontal bar containing three circular icons and their corresponding text labels. From left to right: an envelope icon with the text 'CREATE A SUPPORT CASE', a speech bubble icon with the text 'CHAT WITH A REPRESENTATIVE', and a telephone handset icon with the text 'CALL A SUPPORT REPRESENTATIVE'.

Tip: When contacting Support with a technical issue, please have both the program's log files and the store you are having difficulty with ready.

Compress the program's log files in a .zip file and attach it to your support request.

If the log archive exceeds 10MB, you may not be able to send them as an email attachment. In that case, upload instructions will be provided to you upon request.

Appendix: Understanding Hyper-V backups on a vault

This appendix provides information about how Hyper-V VM backups are stored on a vault.

Note: This information is provided for vault administrators. It might not be relevant for customers who back up Hyper-V VMs to vaults hosted by a service provider.

When you run a Hyper-V Agent backup job, each VM in the job is backed up as a separate job (task) on the vault. This differs from traditional Agent jobs, where each backup job is associated with a single task on the vault.

A task is created on the vault for a VM as soon as a backup job that includes the VM is created. That is, a task for a VM is created on the vault before the VM is backed up.

Beginning in Portal 8.89, the Virtual Machines tab for a Hyper-V Agent on the Computers page shows the vault task name for each protected Hyper-V VM. In previous Portal versions, the vault task name for a Hyper-V VM appeared in a tooltip if you pointed to the VM name.

If a protected VM has been deleted from the Hyper-V environment, and is no longer included in a backup job, you can still see the VM in Portal and restore the VM from the vault.

Determine the name of a VM's task on the vault

Each VM in a Hyper-V backup job is backed up as a separate task on the vault, and is automatically assigned a unique task name. To help you find each task on the vault, you can view the task name for each protected Hyper-V VM in Portal.

Beginning in Portal 8.89, the task name for each VM is shown on the Virtual Machines tab for a Hyper-V Agent. In previous Portal versions, the task name appeared in a tooltip if you pointed to the VM name.

Note: To determine the vault where a particular VM backup is saved, check the VM's backup history. The vault IP address for a safeset appears in the **Location** field on the Backup History tab. See [View a Hyper-V VM's backup history and logs](#). To determine the Account, Username, and the Agent Host name for the backup on the vault, see information in the Vault Settings dialog box. See [Add vault settings](#).

To determine the name of a VM's task on the vault:

1. In Portal, on the navigation bar, click **Computers**.

A grid lists available computers.

2. Find the Hyper-V environment with the protected VM, and expand the environment view by clicking the row.

3. Click the **Virtual Machines** tab.

The Virtual Machines tab shows all protected VMs in the Hyper-V cluster or standalone host.

4. In the Current Inventory/Protected Inventory filter, click **Protected Inventory**.

The Virtual Machines tab shows VMs that have been backed up and can be restored. The VM ID on Vault column shows the vault task name for each VM.