



# Carbonite Server Backup Hyper-V Agent 8.6

User Guide



© 2017 Carbonite, Inc. All rights reserved.

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, Inc.  
Two Avenue de Lafayette  
Boston, MA 02111  
[www.carbonite.com](http://www.carbonite.com)

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

Acknowledgements: Two encryption methods, DES and TripleDES, include cryptographic software written by Eric Young. The Windows versions of these algorithms also include software written by Tim Hudson. Bruce Schneier designed Blowfish encryption.

“Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are Copyright 2001-2006 Robert A. van Engelen, Genivia Inc. All Rights Reserved. THE SOFTWARE IN THIS PRODUCT WAS IN PART PROVIDED BY GENIVIA INC AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.”

The Carbonite Server Backup Agent, Carbonite Server Backup CentralControl, and Carbonite Server Backup Director applications have the encryption option of AES (Advanced Encryption Standard). Advanced Encryption Standard algorithm (named Rijndael, pronounced “Rain Doll”) was developed by cryptographers Dr. Joan Daemen and Dr. Vincent Rijmen. This algorithm was chosen by the National Institute of Standards and Technology (NIST) of the U.S. Department of Commerce to be the new Federal Information Processing Standard (FIPS).

The Carbonite Server Backup Agents and Carbonite Server Backup Director applications also have the added security feature of an over the wire encryption method.

## Document History

Version	Date	Description
1	December 2017	Initial guide provided for Hyper-V Agent 8.6.

## Contents

<b>1</b>	<b>Introduction to the Hyper-V Agent</b> .....	<b>5</b>
1.1	Hyper-V Agent components .....	5
<b>2</b>	<b>Prepare for a Hyper-V Agent deployment</b> .....	<b>7</b>
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	Best practices in a protected Hyper-V environment.....	8
<b>3</b>	<b>Install and upgrade the Hyper-V Agent</b> .....	<b>10</b>
3.1	Install the Hyper-V Agent Management service .....	10
3.2	Install the Hyper-V Agent Host service.....	11
3.3	Upgrade the Hyper-V Agent .....	12
<b>4</b>	<b>Configure the Hyper-V Agent</b> .....	<b>15</b>
4.1	Change credentials or the network address for accessing Hyper-V .....	17
4.2	Undelete Hyper-V environments .....	17
4.3	Add vault settings.....	18
4.4	Add a description .....	20
4.5	Add retention types.....	20
4.6	Set up email notifications for a computer.....	22
4.7	Configure bandwidth throttling .....	23
<b>5</b>	<b>Add and run Hyper-V backup jobs</b> .....	<b>25</b>
5.1	Best practices for backing up Hyper-V VMs .....	25
5.2	Best practices for seeding Hyper-V VM backups .....	26
5.3	Add a Hyper-V backup job.....	27
5.4	Add a Hyper-V backup job by selecting VMs.....	28
5.5	Edit a Hyper-V backup job .....	31
5.6	Add or edit a schedule for a Hyper-V backup job .....	33
5.7	Delete a backup job.....	36
5.8	Disable or enable all scheduled backup jobs.....	36

5.9	Run an ad-hoc backup .....	37
<b>6</b>	<b>Restore Hyper-V VMs .....</b>	<b>39</b>
<b>7</b>	<b>Recover jobs and settings from an offline Hyper-V Agent .....</b>	<b>43</b>
7.1	Hyper-V disaster recovery .....	46
<b>8</b>	<b>Monitor computers and processes .....</b>	<b>48</b>
8.1	View computer and job status information .....	48
8.2	View an unconfigured computer's logs .....	49
8.3	View current process information for a job .....	51
8.4	View a job's process logs and safeset information .....	53
8.5	View and export recent backup statuses .....	55
8.6	View a Hyper-V VM's backup history and logs .....	56
8.7	Hyper-V Agent logs and configuration files .....	58
<b>9</b>	<b>Understanding and troubleshooting Hyper-V processes .....</b>	<b>59</b>
<b>10</b>	<b>Carbonite Server Backup Support .....</b>	<b>60</b>
10.1	Contacting Carbonite .....	60
<b>Appendix A:</b>	<b>Alternate Hyper-V Agent deployments .....</b>	<b>61</b>
A1.	Alternate deployment for protecting a Hyper-V cluster .....	61
A2.	Alternate deployment for a protecting a Hyper-V standalone host .....	61
<b>Appendix B:</b>	<b>Install, upgrade and uninstall the Hyper-V Agent .....</b>	<b>63</b>
B1.	Install the Hyper-V Agent Management service in silent mode .....	63
B2.	Install the Hyper-V Agent Host service in silent mode .....	64
B3.	Upgrade the Hyper-V Agent Management service in silent mode .....	66
B4.	Upgrade the Hyper-V Agent Host service in silent mode .....	66
B5.	Uninstall the Hyper-V Agent Management service .....	67
B6.	Uninstall the Management service in silent mode .....	67
B7.	Uninstall the Hyper-V Agent Host service .....	68
B8.	Uninstall the Host service in silent mode .....	69
<b>Appendix C:</b>	<b>Understanding Hyper-V backups on a vault .....</b>	<b>70</b>
C1.	Determine the name of a VM's task on the vault .....	70

# 1 Introduction to the Hyper-V Agent

Carbonite Server Backup Agent for Hyper-V provides data protection for Microsoft Hyper-V environments, without requiring Agent software to be installed on individual virtual machines (VMs). The Agent protects VMs in standalone and clustered Hyper-V environments, with support for application-consistent backups.

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.

Each VM is independent from the job in which it is backed up. When restoring a VM, you do not need to remember which backup job it was in. You can restore a protected VM even if its backup job has been deleted.

The Hyper-V Agent is closely integrated with Carbonite Server Backup 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 Carbonite or installed on-premises.

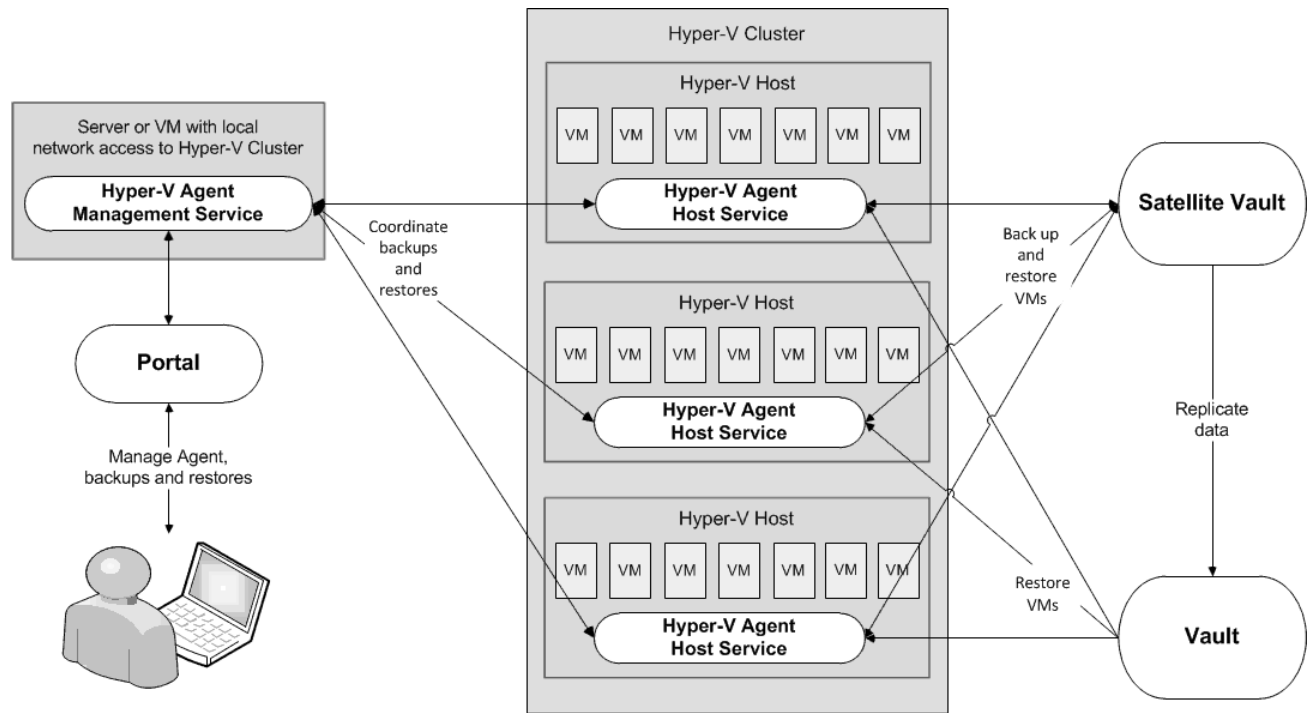
## 1.1 Hyper-V Agent components

The Hyper-V Agent consists of two components:

- **Hyper-V Agent Management service.** The Management service is a central management component that communicates with Carbonite Server Backup 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.

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](#).

### 2.1 Portal for managing a Hyper-V Agent

The Hyper-V Agent is managed using Carbonite Server Backup Portal.

*Note:* You cannot manage the Hyper-V Agent using 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 Carbonite, 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 .

The data can then be replicated to the Carbonite cloud to ensure offsite protection in the case of a disaster.

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

### 2.3 Recommended deployment for protecting a Hyper-V standalone host

To protect a standalone Hyper-V host, Carbonite recommends installing both the Management service and Host service on the standalone host.

If you want to minimize performance impact in the environment, or do not want to open the virtualized environment for Portal or vault access, see the alternate deployment method described in [Alternate deployment for a protecting a Hyper-V standalone host](#).

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

*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, Carbonite recommends the following:

- Install the Management service on a VM in the cluster, and enable High Availability on the VM.
- 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.

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

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

*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.

## 2.5 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 the Hyper-V Agent.

### Enable CSV Cache

In a failover cluster, enabling the CSV cache might improve Hyper-V Agent backup performance. Microsoft recommends enabling the CSV cache for read-intensive workloads. See “Use Cluster Shared Volumes in a Failover Cluster” (<http://technet.microsoft.com/en-us/library/jj612868.aspx>) and “How to Enable CSV Cache” (<http://blogs.msdn.com/b/clustering/archive/2012/03/22/10286676.aspx>).

### Clean up snapshots and checkpoints before backups

The Hyper-V Agent backs up and restores user-level snapshots or checkpoints with VMs, which can take a significant amount of time.

*Note:* “Snapshots” in Windows Server 2012 are the same as “checkpoints” in later Windows Server versions.

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. See “Hyper-V: Avoid using snapshots on a virtual machine that runs a server workload in a production environment” ([http://technet.microsoft.com/en-us/library/ee941140\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/ee941140(v=ws.10).aspx)).



## Use fixed-size VHDXs or VHDs

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

Consistent with Microsoft best practices, we recommend not using dynamically expanding VHDXs or VHDs in a production environment. See “Hyper-V: VHD-format dynamic virtual hard disks are not recommended for virtual machines that run server workloads in a production environment”

(<http://social.technet.microsoft.com/wiki/contents/articles/13078.hyper-v-vhd-format-dynamic-virtual-hard-disks-are-not-recommended-for-virtual-machines-that-run-server-workloads-in-a-production-environment.aspx>).

## Avoid using VMs with limited or no backup support

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). See “Planning for Backup”

([http://technet.microsoft.com/en-us/library/dd252619\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd252619(WS.10).aspx)).

During a backup, the Hyper-V Agent skips VMs that contain mixed storage or share virtual hard disks. See [Understanding and troubleshooting Hyper-V processes](#).

The Hyper-V Agent cannot back up a VM with 50 or more checkpoints. Microsoft specifies a maximum of 50 checkpoints for a VM. See “Maximums for virtual machines” (<https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/plan/plan-hyper-v-scalability-in-windows-server>).

## 3 Install and upgrade the Hyper-V Agent

To protect a Microsoft Hyper-V environment, you must install two Hyper-V Agent components:

- Hyper-V Agent Management service. See [Install the Hyper-V Agent Management service](#).
- 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.

You can then configure the Hyper-V Agent to communicate with the Hyper-V environment, and add a vault connection. See [Configure the Hyper-V Agent](#).

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

### 3.1 Install the Hyper-V Agent Management service

The Hyper-V Agent Management service is installed on a server that has local network access to a protected Hyper-V environment. See [Recommended deployment for protecting a Hyper-V cluster](#) and [Recommended deployment for protecting a Hyper-V standalone host](#).

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.

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

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. On the **Welcome** page, click **Next**.
3. On the **License Agreement** page, read the license agreement. Click **I accept the terms in the license agreement**, and then click **Next**.
4. On the **Destination Folder** page, do one of the following:
  - To install the Management service in the default location (C:\Program Files\Carbonite Server Backup\Hyper-V Agent Management), 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**.

5. On the **Register Carbonite Server Backup Hyper-V Agent Management with Portal** page, specify the following information:
  - In the **Network Address** box, type the host name or IP address of the Portal for managing the Hyper-V Agent. Specifying the host name of the Portal is recommended. This will allow DNS to handle IP address changes.
  - 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.
6. Click **Next**.
7. 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.
8. On the **Ready to Install the Program** page, click **Install**.

## 3.2 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 [Recommended deployment for protecting a Hyper-V cluster](#) and [Recommended deployment for protecting a Hyper-V standalone host](#).

Before you can install a Host service, the Hyper-V Agent Management service must be installed on a server with local network access to the Hyper-V environment. During the installation, the Host service must be able to establish connection with the Management service. Ensure that there is local network connectivity to the Management service and that the correct port is open.

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

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

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. 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 Host service in the default location (C:\Program Files\Carbonite Server Backup\... ), click **Next**.
  - To specify another installation location, click **Change**. In the **Change Current Destination Folder** window, browse to the new installation location, or enter a folder in the **Folder name** box. Click **OK**. On the **Destination Folder** page, click **Next**.
6. On the **Connect with Hyper-V Agent Management service** page, in the **Network Address** box, enter the host name or IP address of the Hyper-V Agent Management service that will assign work to the Host service. Specifying the host name of the Management service is recommended. This will allow DNS to handle IP address changes.
7. 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.
8. Click **Next**.
9. Click **Install**.
10. On the **Wizard Completed** page, click **Finish**.

### 3.3 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 services must be upgraded to the same version. Earlier service versions cannot be used with later service versions.

During a cluster rolling upgrade to Windows Server 2016, the Hyper-V Agent can continue to back up VMs. To ensure that backups continue during a rolling upgrade, upgrade the Hyper-V Management service first, upgrade the Hyper-V Agent Host service on each cluster node, and then upgrade each cluster node to Windows Server 2016.

You can also move to a new Agent version when recovering a protected Hyper-V environment after a disaster or when migrating VMs to a new environment. See [Upgrade the Hyper-V Agent in a new environment](#).

#### 3.3.1 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.

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

*Note:* 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**.
3. In the installation wizard, click **Next**.
4. On the Installation Completed page, click **Finish**.

### 3.3.2 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](#).

*Note:* 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**.
3. In the installation wizard, click **Next**.
4. On the Installation Completed page, click **Finish**.

### 3.3.3 Upgrade the Hyper-V Agent in a new environment

When recovering a protected Hyper-V environment after a disaster or when migrating VMs to a new environment, you can move to a new Agent version and recover Agent jobs and settings from the previous Agent version. For example, you can move from Hyper-V Agent 7.30 in a Windows Server 2012 environment to Hyper-V Agent 7.40 in a Windows Server 2012 R2 environment, recover backup jobs, and back up VMs in the new environment using the recovered jobs.

To upgrade the Hyper-V Agent when migrating to a new environment:

1. Create a new Hyper-V environment.
2. If possible, migrate VMs from the original environment to the new Hyper-V environment.
3. Install the new Hyper-V Agent version in the new Hyper-V environment. See [Install the Hyper-V Agent Management service](#) and [Install the Hyper-V Agent Host service](#).
4. Recover jobs and settings from the earlier Hyper-V Agent version. See [Recover jobs and settings from an offline Hyper-V Agent](#).

5. If the original protected Hyper-V environment is lost or unavailable, or VMs cannot be migrated to the new environment, restore protected VMs to the new Hyper-V environment. See [Restore Hyper-V VMs](#).
6. Enable all scheduled jobs for the new Hyper-V Agent. See [Disable or enable all scheduled backup jobs](#).

## 4 Configure the Hyper-V Agent

After the Hyper-V Agent Management service is installed and registered with Carbonite Server Backup Portal, you must configure the Agent by doing the following:

- Provide credentials for authenticating with the Hyper-V environment that you want to protect. The user should be an Active Directory domain user with administrative rights to the Hyper-V cluster or standalone host.
- 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.

*Note:* In earlier Portal versions, you could specify whether data is encrypted using AES encryption when it is transmitted to and from the vault. Over-the-wire encryption is now automatically enabled when you add vault settings or save existing vault settings.

You can also change Hyper-V credentials and add vault settings after the initial configuration. See [Change credentials or the network address for accessing Hyper-V](#) and [Add vault settings](#).

Optionally, you can do the following:

- 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 a computer](#).
- Specify the amount of bandwidth consumed by backups. See [Configure bandwidth throttling](#).

To configure the 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.

The **Configuration mode selection** section appears.

3. Select **Configure a new Hyper-V Agent**, and then click **Continue**.

*Note:* The **Recover a previous Hyper-V Agent** option will also appear if offline Hyper-V Agents are available. This option is used to recover Hyper-V Agent configuration and all backup jobs from an offline 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 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.
  - 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.
  - In the **Username** box, type the domain 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*.
  - 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 message appears. Click **Okay**.
6. Click **Continue**.
7. In the **Vault Configuration** section, click **Configure Vault**.

*Note:* 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 IP 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.

For more information about policies, see the Portal online help.
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.
  - **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. If the Agent cannot connect to the vault successfully in the specified number of tries, the backup or restore fails.



11. Click **Save**.

The name of the Hyper-V cluster or standalone host now appears on the Computers page in Portal instead of the Management service computer name.

## 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**. Find the 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.
  - 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.
  - In the **Username** box, type the domain 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*.
  - 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**.

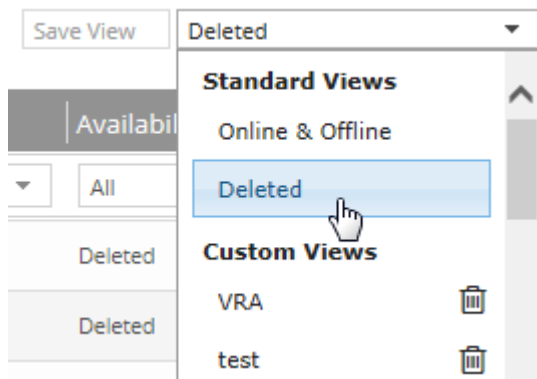
## 4.2 Undelete Hyper-V environments

You can view Hyper-V environments that have been deleted from Portal, and “undelete” deleted Hyper-V environments.

If a Hyper-V environment has been deleted from Portal, you must undelete the Hyper-V environment before you can recover jobs and settings from the environment. 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**.

### 4.3 Add vault settings

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

When adding vault settings for an Agent, 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 an Agent, 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 an Agent, 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 the Hyper-V Agent](#).

In previous Portal versions, you could specify whether data is encrypted using AES encryption when it is transmitted to and from the vault. Over-the-wire encryption is now automatically enabled when you add vault settings or save existing vault settings.

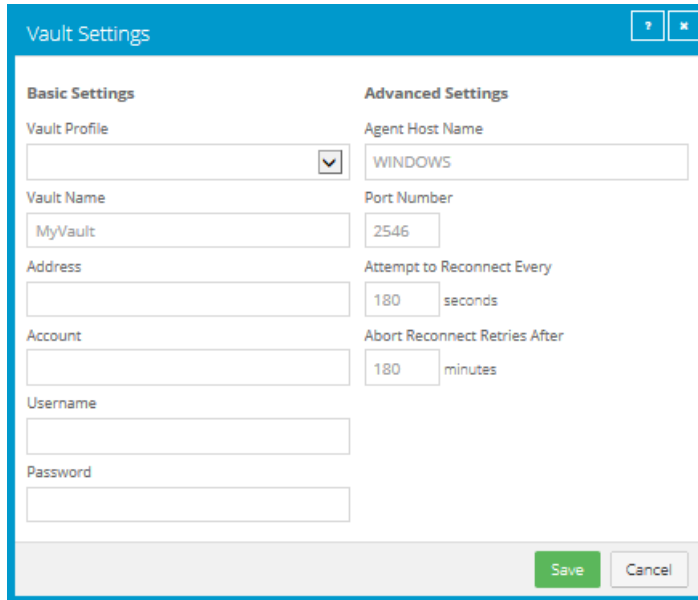
When an E2 appliance reports a new IP address, the IP address is updated in Portal vault settings for Agents that are registered to the E2, and in the E2 vault profile. Agent versions 8.10 and later contact Portal to check for vault IP address changes. If a Super user or Admin user changes the name of an E2 vault profile, the name is updated automatically in vault settings for Agents that are registered to the E2.

To add vault settings:

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

2. Find the Agent for which you want to add vault settings, and click the computer row to expand its view.
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 IP 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.

- **Abort Reconnect Retries After.** Specifies the number of times the Agent tries to reconnect to the vault, if the vault becomes unavailable during a backup or restore. If the Agent cannot connect to the vault successfully in the specified number of tries, the backup or restore fails.

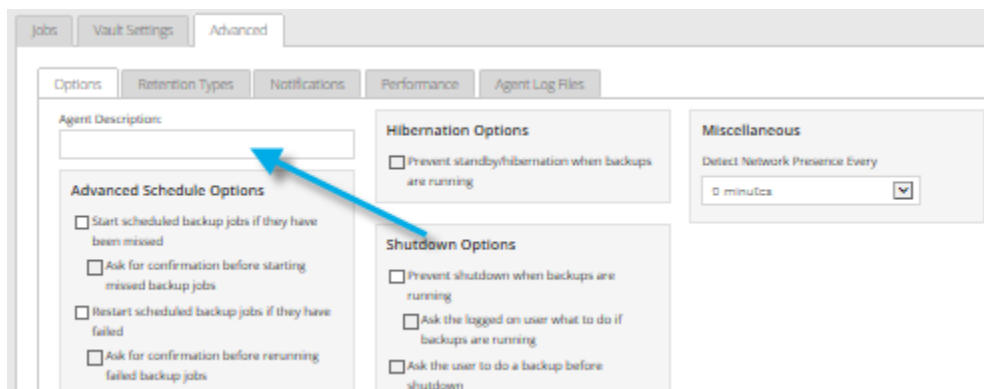
6. Click **Save**.

## 4.4 Add a description

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

To add a description:

1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to add a description, and click the row to expand its view.
3. On the **Advanced** tab, click the **Options** tab.
4. In the **Agent Description** box, enter a description for the Agent.



5. Click **Save**.

## 4.5 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 an Agent where a policy is not assigned.

If a policy is assigned to an Agent, retention types cannot be added or modified on the Computers page. Instead, retention types can only be added or modified in the policy. See the Portal online help.

To add a retention type:

1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to add a retention type, and click the row to expand its view.
3. On the **Advanced** tab, click the **Retention Types** tab.

If a policy is assigned to the Agent, 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>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**.

## 4.6 Set up email notifications for a computer

To make it easier to monitor backups, users can receive emails when backups finish or fail.

Admin users and regular users in Portal can set up email notifications for a computer.

To set up email notifications for a computer:

1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to configure email notifications, and click the 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 Agent, you cannot change values on the **Notifications** tab. Instead, notifications can only be modified in the policy.

The screenshot shows the configuration interface for an agent, with the 'Advanced' tab selected. Under the 'Advanced' tab, the 'Notifications' sub-tab is active. At the top, there are three checkboxes: 'On Successful Completion', 'On Failure', and 'On Error'. Below these are two main sections: 'SMTP Settings' and 'SMTP Credentials (if required)'. The 'SMTP Settings' section contains four input fields: 'Email "From" Address:', 'Outgoing Mail Server (SMTP):', 'Recipient Address(es):', and 'Outgoing Server Port (SMTP):'. The 'SMTP Credentials (if required)' section contains three input fields: 'User Name:', 'Password:', and 'Domain:'.

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.

4. Click **Save**.

## 4.7 Configure bandwidth throttling

Bandwidth throttling settings specify the amount of bandwidth consumed by an Agent for backups. For example, you might want to restrict the amount of bandwidth used for daytime backups so that online users are not affected, and allow unlimited bandwidth usage at night so that scheduled backups run as fast as possible.

For the Hyper-V Agent, bandwidth throttling is applied at the Host level. If three VMs are being backed up on a node, each gets 1/3 of the specified maximum bandwidth on the node. 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 all 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.

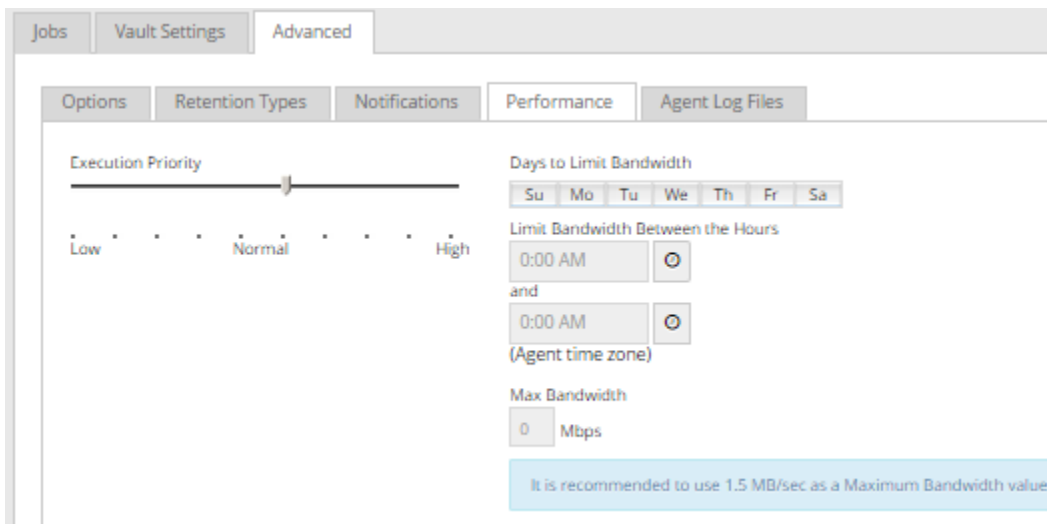
If you edit an Agent's bandwidth settings while a backup is running, the new Agent settings do not affect the backup that is running. Bandwidth settings are applied when a backup starts, and are not applied to backups that are already running.

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. See the Portal online help.

To configure bandwidth throttling:

1. On the navigation bar, click **Computers**.
2. Find the Agent for which you want to configure bandwidth throttling, and click the row to expand its view.
3. Click the **Advanced** tab, click the **Performance** tab, and then edit the bandwidth settings.

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





## 5 Add and run Hyper-V backup jobs

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.

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](#) or [Add a Hyper-V backup job by selecting VMs](#).

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 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.
- 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.

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 Hyper-V backup jobs.

#### **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.

#### **Include VMs on the same CSV in the same backup job**

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.

## Avoid 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 is directed to a different vault.
- VM backups in a job reseed if the job’s encryption password changes.
- If a VM is backed up as part of one backup job, and is then moved to a job with a different encryption password, the amount of data sent to the vault is equivalent to a seed backup. If a VM is moved to a different backup job with the same encryption password, the VM backup does not reseed.
- After storage migration, snapshot (AVHD) files reseed. Hard disk VHD(x) files do not reseed after storage migration.

## 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 an appliance to provide fast, local vault access. If you do not use an appliance, consider using a temporary vault to seed Hyper-V backups locally. The data can then be imported into the Carbonite cloud.

### 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, Carbonite recommends 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](#) or [Add a Hyper-V backup job by selecting VMs](#).

- 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](#).
- 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 Add a Hyper-V backup job

To add a Hyper-V backup job:

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

The Computers page shows registered computers and environments.

- Click the Hyper-V environment row. 
- Click the **Jobs** tab.
- In the **Job Tasks** 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.

- 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.

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

- To find one or more VMs in the **Unprotected** or **Protected** 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.

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

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

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 **Create Job**.

## 5.4 Add a Hyper-V backup job by selecting VMs

When creating a new Hyper-V backup job, you can select VMs on the Virtual Machines tab to include in the job. You can only select VMs that appear on the Hyper-V Virtual Machines tab at the same time.


*Note:* 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 add a Hyper-V backup job without selecting VMs on the Virtual Machines tab. See [Add a Hyper-V backup job](#).

To add a Hyper-V backup job by selecting VMs:

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

The Computers page shows registered computers and environments.

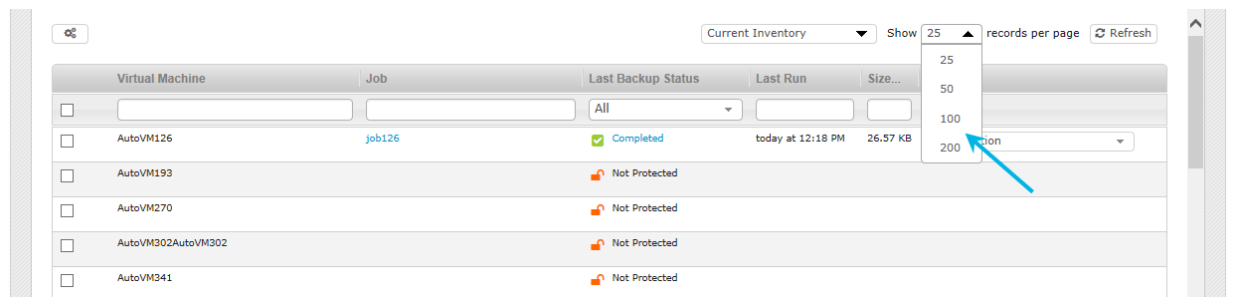
2. Click the Hyper-V environment row. 

3. 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.

4. By default, the **Virtual Machines** tab shows 25 VMs at a time. 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.

*Note:* 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.



5. 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.

*Note:* If you select a VM that is already included in another job, it will not be added to the new job.

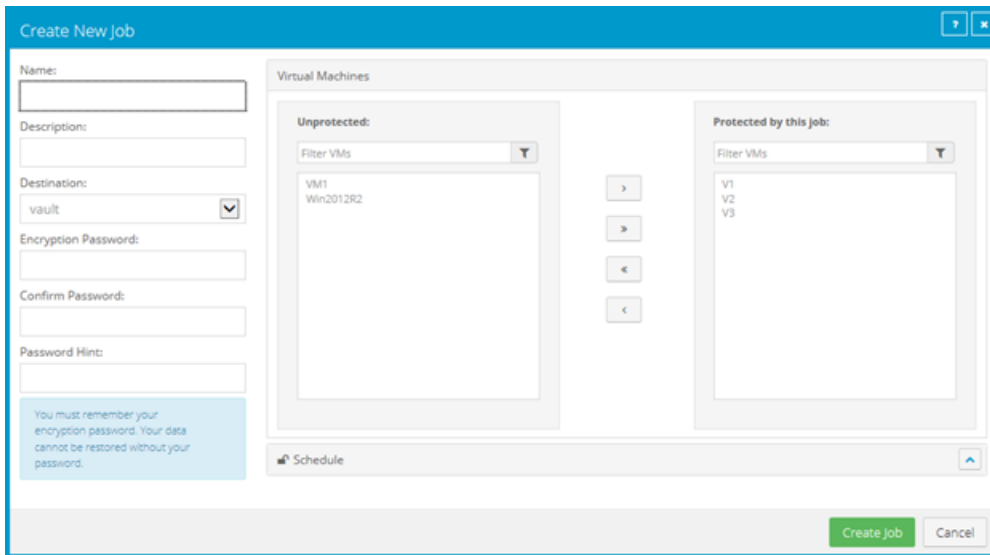
6. Click **Create Hyper-V Job**. 

7. 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.



8. 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.

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

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

9. 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](#).
10. 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
- Schedules and retention types

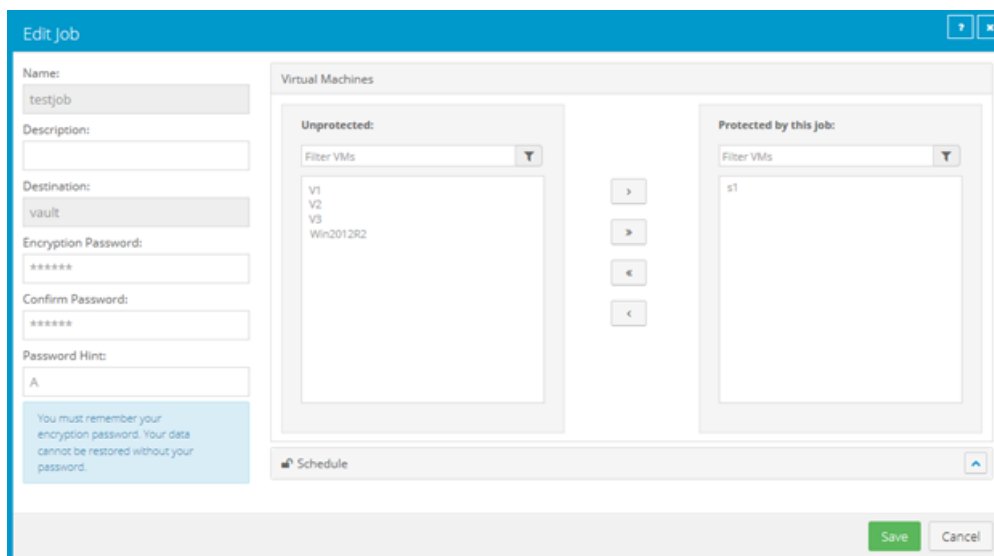
You cannot change a backup job's name or vault connection.



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. In the **Edit Job** dialog box, change the job description, encryption password, or password hint, if desired.

VMs in the backup job appear in the **Protected by this job** box.



5. 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** 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.

- 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, 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.

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. See [Add a Hyper-V backup job by selecting VMs](#).

6. 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](#).
7. 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 job is scheduled to run at slightly different times, the Agent attempts to run the job according to each schedule. For example, if a job is scheduled to run at 11 PM by one schedule and 11:01 PM by another schedule, the Agent will attempt to run the job twice. Try to avoid overlapping schedules; problems can occur if a job is scheduled to run twice in a short period of time.

Retention	Schedule	Enable	Priority
Monthly	11:00 PM Last	<input checked="" type="checkbox"/>	High
Daily	11:00 PM Su,Mo,Tu,We,Th,Fr	<input checked="" type="checkbox"/>	Low

*Note:* You cannot defer scheduled Hyper-V backups. Hyper-V Agent backups can only be deferred when they are run manually (ad hoc). See [Run an ad-hoc backup](#).

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.

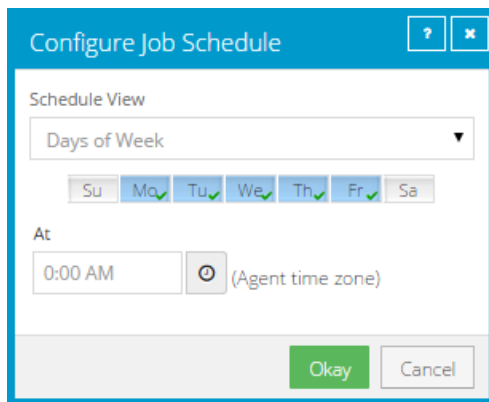
A retention type specifies the number of days a backup is kept on the vault, how many copies of a backup are stored online, and how long backup data is stored offline.

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

The **Configure Job Schedule** dialog box opens.

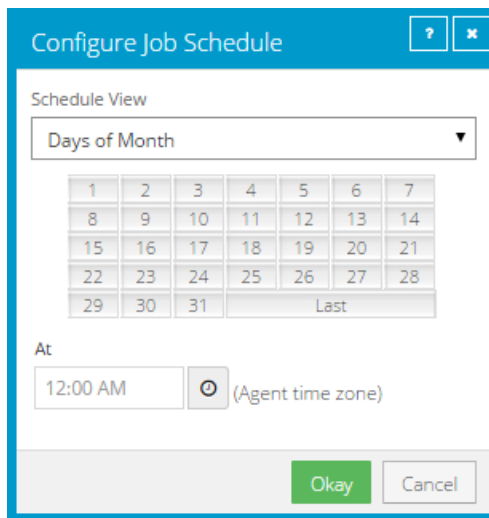
5. In the **Configure Job Schedule** dialog box, do one of the following:

- 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**.



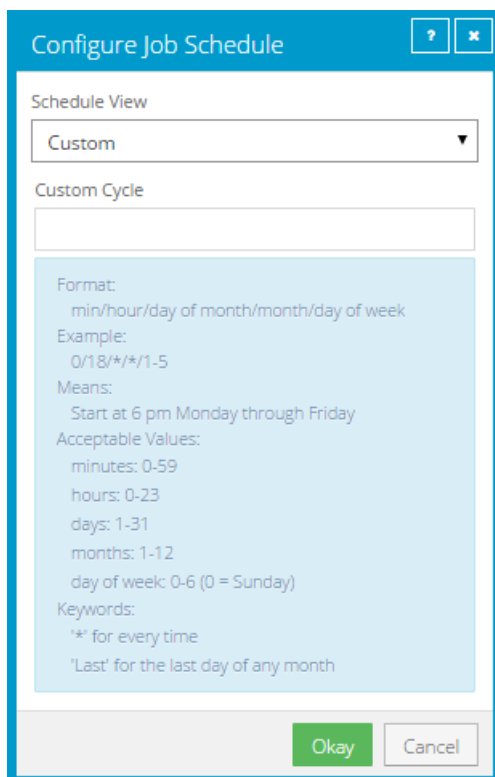
The screenshot shows the 'Configure Job Schedule' dialog box with the 'Schedule View' dropdown set to 'Days of Week'. Below the dropdown, there are seven buttons representing the days of the week: Su, Mo, Tu, We, Th, Fr, and Sa. The Mo, Tu, We, Th, and Fr buttons have a small green checkmark in the bottom right corner, indicating they are selected. Below the day buttons is the 'At' field, which contains '0:00 AM' and a clock icon, followed by the text '(Agent time zone)'. At the bottom of the dialog box are two buttons: 'Okay' (highlighted in green) and 'Cancel'.

- 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**.



The screenshot shows the 'Configure Job Schedule' dialog box with the 'Schedule View' dropdown set to 'Days of Month'. Below the dropdown is a calendar grid with dates from 1 to 31, and a 'Last' button. The 'At' field contains '12:00 AM' and a clock icon, followed by the text '(Agent time zone)'. At the bottom of the dialog box are two buttons: 'Okay' (highlighted in green) and 'Cancel'.

- 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 image shows a dialog box titled "Configure Job Schedule" with a blue header bar containing a question mark icon and a close button. The dialog is divided into several sections:

- Schedule View:** A dropdown menu currently showing "Custom".
- Custom Cycle:** An empty text input field.
- Format:** A light blue box containing the following text:
  - min/hour/day of month/month/day of week
  - Example: 0/18/\*/\*/1-5
  - Means: Start at 6 pm Monday through Friday
  - Acceptable Values:
    - minutes: 0-59
    - hours: 0-23
    - days: 1-31
    - months: 1-12
    - day of week: 0-6 (0 = Sunday)
  - Keywords:
    - \* \* for every time
    - 'Last' for the last day of any month

At the bottom of the dialog, there are two buttons: a green "Okay" button and a grey "Cancel" button.

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 Delete a backup job

*Note:*

To delete a backup job:

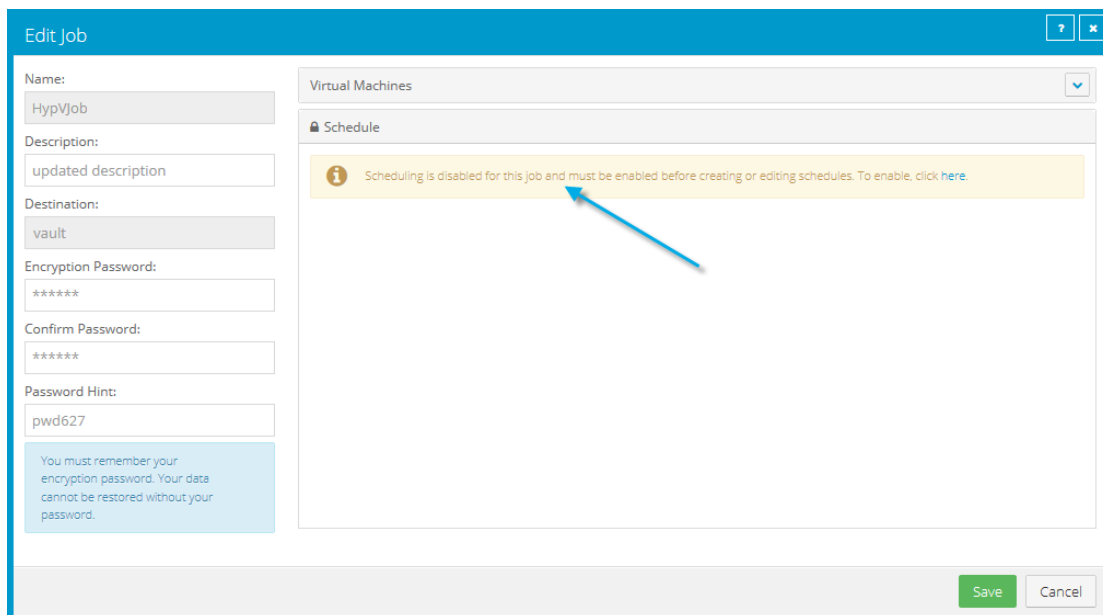
1. On the navigation bar, click **Computers**.  
The Computers page shows registered computers.
2. Find the Agent with the job that you want to delete, and expand its view by clicking its row.
3. Click the **Jobs** tab.
4. In the **Select Action** menu of the job that you want to delete, click **Delete Job**.
5. In the confirmation dialog box, click **Delete**.

## 5.8 Disable or enable all scheduled backup jobs

Admin users can disable or enable all scheduled backup jobs for a Hyper-V environment.

*Note:* You can also disable or enable a specific schedule for a backup job. See Add or edit a schedule for a Hyper-V backup job.

When you disable all scheduled jobs for a protected environment, backup jobs do not run according to any schedules. 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 (as shown in the following screenshot).



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 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**.

## 5.9 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 VMs

You can restore one or more virtual machines (VMs) from a Hyper-V backup. In a single request, you can restore VMs that were backed up using multiple Hyper-V backup jobs, even if each job has a different encryption password. You can restore a protected VM even if the original VM no longer exists in the Hyper-V environment and its backup job no longer exists.

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

If you restore VMs from a Windows 2012 environment to a Windows 2012 R2 environment, the restored VMs are restored as Generation 1 VMs. You cannot restore VMs from a Windows 2016 or 2012 R2 host to a Windows 2012 environment.

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.

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.

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.

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 nodes 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.

In the event of a disaster, a temporary vault can be used to provide local vault access for restoring VMs that are backed up to the Carbonite cloud.

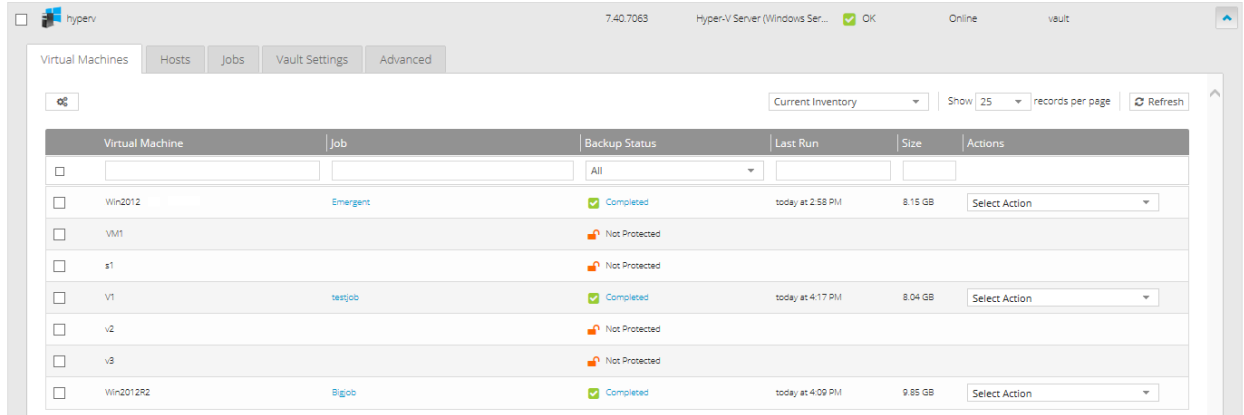
To restore Hyper-V VMs:

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

A grid lists available computers.

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


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



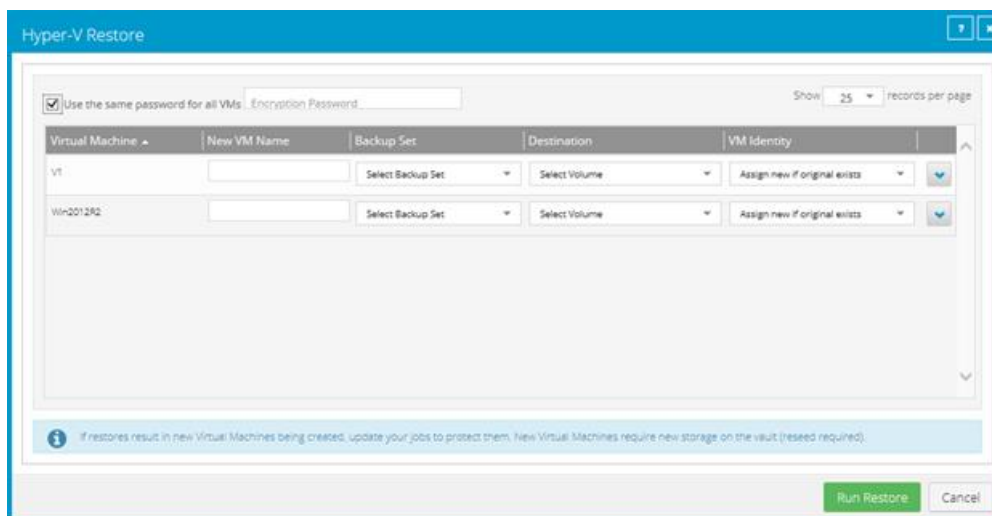
- In the Current Inventory/Protected Inventory filter, click **Protected Inventory**.

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

- 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**. 

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




- Do one of the following:

- If you are restoring one VM, enter the data encryption password in the **Encryption Password** box.



- 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.

7. For each VM that you are restoring, do the following in the VM row:

- (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**.

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

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 system volumes. System volumes do not appear in the **Destination** list.

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

*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.
- 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**.

8. Click **Run Restore**.

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

If a Hyper-V Agent goes offline because it is lost or unavailable, you can install a new Management service and recover jobs and settings from the offline Agent. You can then enter credentials, run backup jobs from the original Agent, and restore VMs that were protected by the Agent.

You can 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 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 an SMTP password for notifications.

You can recover jobs and settings from an earlier Hyper-V Agent version, or the current version. Recovering jobs from a previous Agent version can be useful when migrating Hyper-V VMs to a new environment. See [Upgrade the Hyper-V Agent](#).

If you recover jobs and settings from an offline Hyper-V Agent in a Windows 2012 environment, and then restore VMs to a Windows 2012 R2 environment, the restored VMs are restored as Generation 1 VMs. You cannot restore VMs from a Windows 2012 R2 environment to a Windows 2012 environment.

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 Carbonite Server Backup 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. On the **Register CHyper-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. See [Install the Hyper-V Agent Management service](#).
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.

- 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.

This agent must be configured before you can continue. The wizard below will help you accomplish this in a few easy steps.

**1** Configuration mode selection

Configure a new Hyper-V Agent

Gives you the option to register an agent with Hyper-V Server.

Recover a previous Hyper-V Agent

Gives you the option to recover a registration and all backup jobs from an offline agent. You can still change all configuration and jobs after recovery.

Continue

- 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 protected Hyper-V environments 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.

This agent must be configured before you can continue. The wizard below will help you accomplish this in a few easy steps.

**1** Configuration mode selection

Configure a new Hyper-V Agent

Gives you the option to register an agent with Hyper-V Server.

Recover a previous Hyper-V Agent

Gives you the option to recover a registration and all backup jobs from an offline agent. You can still change all configuration and jobs after recovery.

Continue

**2** Recover

Select an offline agent to recover from

AC-DEV, (1/31/2014 10:14:34 AM -05:00) ▼

Continue

- From the **Select an offline agent to recover from** list, choose the Hyper-V Agent 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 Hyper-V Agent's jobs and settings. On the **Computers** page, the computer name changes to the name of the protected Hyper-V cluster or standalone host.

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

7. Click **OK**.
8. 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, enter its encryption password, and 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 a computer](#).
13. Click **Save**. In the confirmation message box, click **Okay**.
14. If the Hyper-V Agent Management service now has a new IP address, check whether Agent Host services are communicating with the Management service. If not, reinstall the Host service on each Hyper-V cluster host or standalone node. See [Install the Hyper-V Agent 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

The following table outlines the process of recovering a protected Hyper-V environment when one or more of the following components are lost:

- 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
- Carbonite Server Backup Portal

**IMPORTANT:** Configuration data, vault, and job information for the Hyper-V Agent is saved in the Carbonite Server Backup 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	
✓			Install the Hyper-V Agent Management service, and recover configuration information and jobs from the offline Hyper-V Agent. See <a href="#">Recover jobs and settings from an offline Hyper-V Agent</a> .
✓	✓		<ol style="list-style-type: none"> <li>1. Rebuild the lost Hyper-V cluster or standalone host (i.e., Windows server with Hyper-V role).</li> <li>2. 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 <b>Cluster Credentials</b> tab. See <a href="#">Recover jobs and settings from an offline Hyper-V Agent</a>.</li> <li>3. Install the Hyper-V Agent Host service on each host, and register it to the Management service. See <a href="#">Install the Hyper-V Agent Host service</a>.</li> <li>4. Restore VMs. See <a href="#">Restore Hyper-V VMs</a>.</li> </ol>

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

## 8 Monitor computers and processes

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

- **Computer page.** The Computer page shows status information for protected environments 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 computers 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](#).
- **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 jobs process logs and safeset information](#) and [View a Hyper-V VMs 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 and export recent backup statuses](#).

### 8.1 View computer and job status information

On the Computer page in Portal, you can view status information for protected environments and their jobs.





To view computer and job status information:


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

The Computers page shows registered Agents.

The **Availability** column indicates whether each Agent 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.
2. Find the Agent for which you want to view logs, and click the row to expand its view.
  3. View the **Jobs** tab.




If a backup or restore is running for a job, an “In Progress” symbol  appears beside the job name, along with the number of processes that are running.



	Name	Job Type	Description
1	AppAware	Image	
2	FilesAndFolders	Local System	







If you click the symbol, the **Process Details** dialog box shows information about running, queued and recently-completed processes for the job. See [View current process information for a job](#).

The **Last Backup Status** column shows the result of the last backup attempt for each job. 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.

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

-  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. For Hyper-V environments, this status can appear when problems were encountered during the backup, but the backups later succeeded.
- 
-  Failed — Indicates that the backup failed and no safeset was created.
-  Cancelled

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


## 8.2 View an unconfigured computer's logs

You can view logs for unconfigured computers. 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 the unconfigured computer, 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.

```


Log Name: BUAgent-LXLOG
25-Nov 06:21:49 AGIIT-I-04314 Agent Version 8.30.7893 Nov 16 2016 14:12:22
25-Nov 06:21:49 AGIIT-I-08103 Executing agent as SYSTEM
25-Nov 06:21:49 AGIIT-I-08199 Agent with Id 216bbd19-cbb7-4176-8dfe-be885ee7ecf7 will connect to server qa.corp.com on port 8086
25-Nov 06:21:49 AGIIT-I-07466 WIII-4 thread started
25-Nov 06:21:49 AGIIT-I-08200 Agent HTTP thread started
25-Nov 06:21:49 AGIIT-I-08200 Agent HTTP thread started
25-Nov 06:21:49 AGIIT-I-08200 Agent HTTP thread started
25-Nov 06:21:50 AGIIT-I-08323 Agent is being redirected to server qa.corp.com on port 8087
25-Nov 06:21:50 AGIIT-I-09400 Agent HTTP binding to 127.0.0.1:8031
25-Nov 06:21:50 AGIIT-I-09400 Agent HTTP binding to :8031
25-Nov 06:21:54 AGIIT-I-07466 WIII-4 thread started
25-Nov 06:21:55 AGIIT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:01 AGIIT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:11 AGIIT-E-08307 Failed to set the Agent status to offline.
25-Nov 06:22:16 AGIIT-I-08914 Agent type set to SERVER
25-Nov 06:22:16 AGIIT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:21 AGIIT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:26 AGIIT-E-07514 Failed to Upload System Info in Notification Thread
25-Nov 06:22:31 AGIIT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:36 AGIIT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:41 AGIIT-E-07477 Failed to Upload Feature Options in Notification Thread
25-Nov 06:22:46 AGIIT-E-07476 Failed to Upload Job Types in Notification Thread


```

### 8.3 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. Process information is typically deleted within an hour after the process ends.

To view current process information for a job:

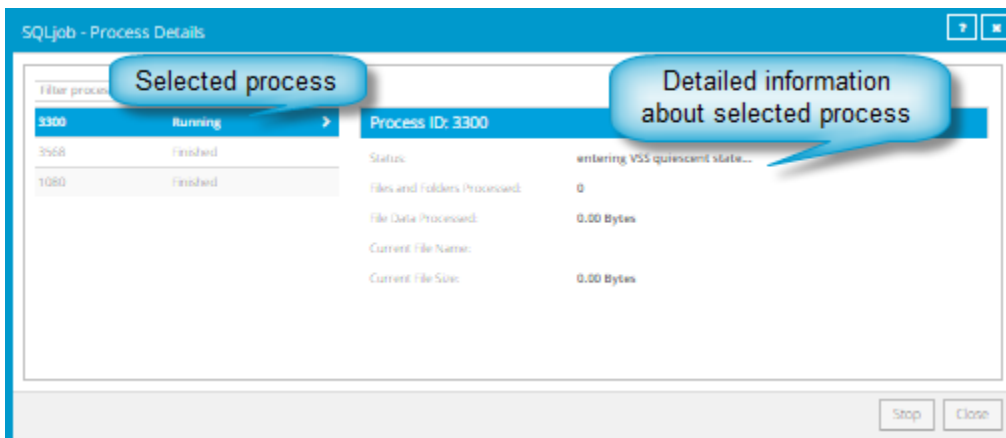
1. Do one of the following:
  - On the Computers page, on the Jobs tab, start a backup, restore or synchronization.
  - On the Computers page, on the Jobs tab, click the “In Progress” symbol  beside the job name.



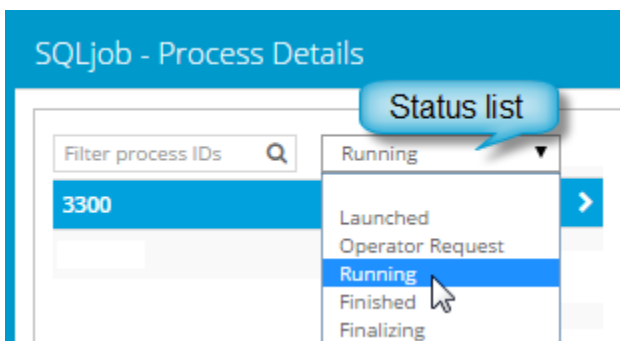
- On the Monitor page, click the “In Progress” symbol  beside the job name.



The **Process Details** dialog box lists 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 on the left side of the dialog box. Detailed information for the process is shown at the right side of the dialog box.
- To show only some processes in the dialog box, do one of the following in the status list:
  - 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**.



## 8.4 View a job's process logs and safeset information

To determine whether a backup or restore 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. For most Agents, one safeset is created by each successful backup.

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

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

The Computers page shows registered Agents.

2. Find the Agent for which you want to view logs, and click the row to expand its view.

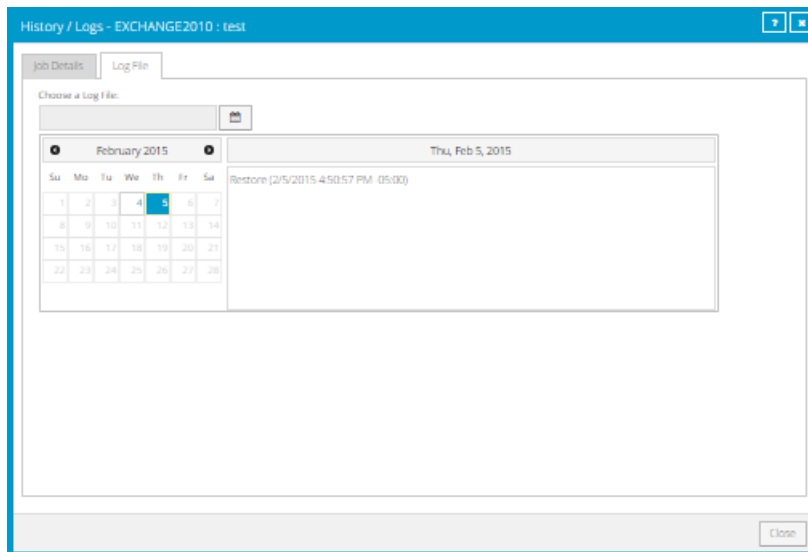
On the **Jobs** tab, the **Last Backup Status** column shows the status of each backup job.


Name	Job Type	Description	Last Backup Status	Last Run	Action
job	Local System		Completed	yesterday at 11:49 PM	Select Action
job2	Local System		Completed	today at 1:32 PM	Select Action

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

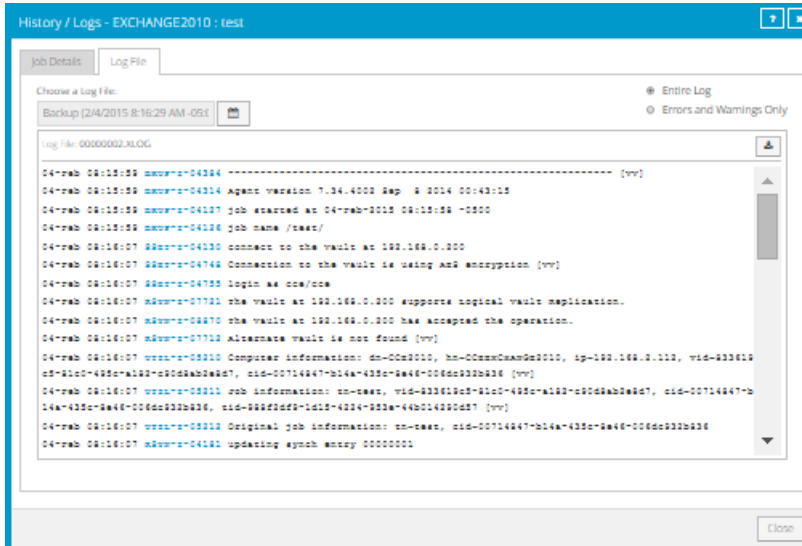
- In the job's **Select Action** menu, click **History / Logs**.
- In the **Last Backup Status** column, click the job status.

The **History / Logs** window lists the most recent backups, restores and synchronizations on the computer.




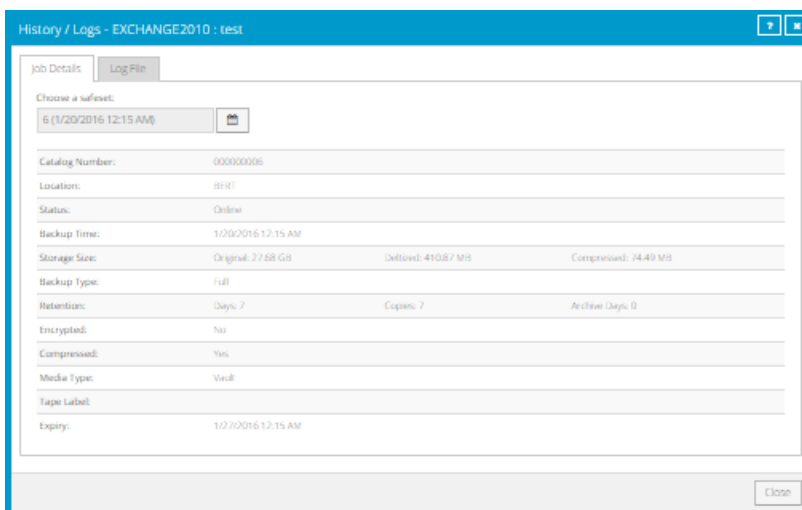
- To view processes for a different day, click the calendar button.  In the calendar that appears, click the date of the log that you want to view. In the list of processes on the selected date, click the process for which you want to view the log.

The **History / Logs** window shows the selected log.



- To only show errors and warnings in the log, click the **Errors and Warnings Only** option at the top right of the window.
- To view safeset information for a particular backup, click the **Job Details** tab. The tab shows safeset information for the job'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.



## 8.5 View and export recent backup statuses

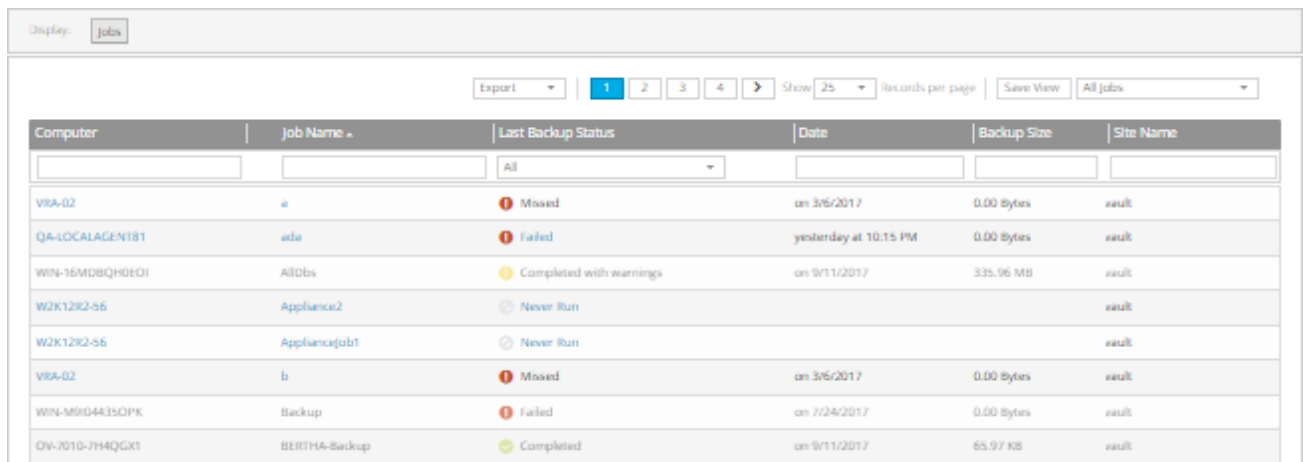
You can view recent backup statuses for computers on the Monitor page in Portal. You can also export the information in comma-separated values (.csv), Microsoft Excel (.xls), or Adobe Acrobat (.pdf) format.

From the Monitor page, you can navigate to related information on the Computers page or in the Logs window.

To view and export recent backup statuses:


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

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



Computer	Job Name	Last Backup Status	Date	Backup Size	Site Name
VRA-02	a	Missed	on 3/6/2017	0.00 Bytes	naul
QA-LOCALAGENT81	ada	Failed	yesterday at 10:15 PM	0.00 Bytes	naul
WIN-16MD8QH0ED1	AllDbs	Completed with warnings	on 9/11/2017	335.96 MB	naul
W2K12R2-56	Appliance2	Never Run			naul
W2K12R2-56	ApplianceJob1	Never Run			naul
VRA-02	b	Missed	on 3/6/2017	0.00 Bytes	naul
WIN-M910443DPK	Backup	Failed	on 7/24/2017	0.00 Bytes	naul
OV-7010-7H4QGK1	BERTHA-Backup	Completed	on 9/11/2017	65.97 KB	naul

2. To change which backup statuses appear on the page, click the views list at the top of the page, and then click the view that you want to apply.
3. To view information for a job or computer on the Computers page, click the name of an online computer or job.
4. To view the job's logs in the History/Logs window, click the job's last backup status.
5. To export backup status information from the 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.

## 8.6 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 Agents.

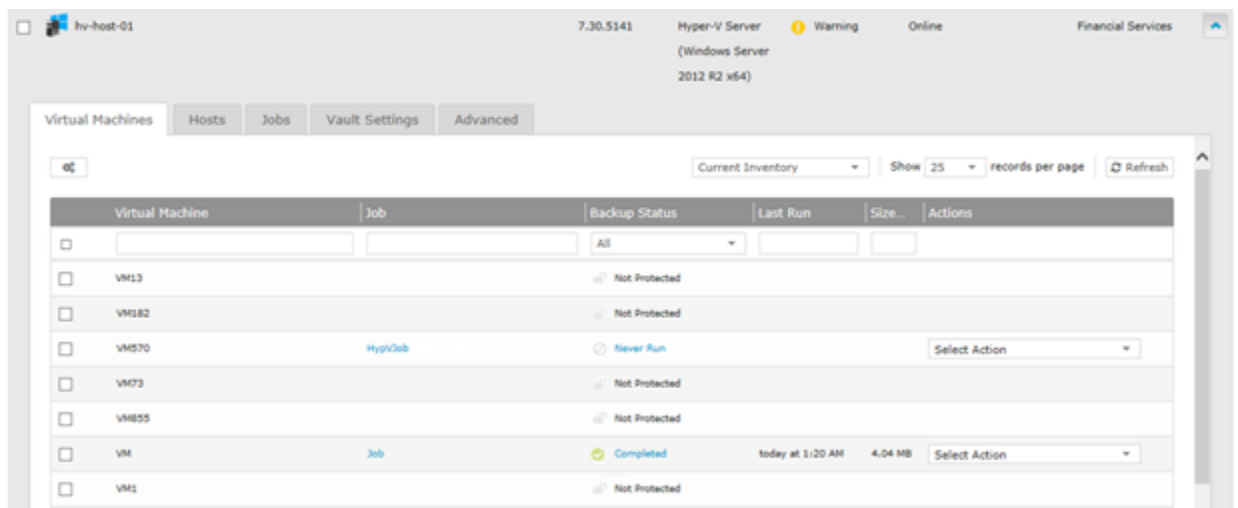
2. Find the Hyper-V Agent for which you want to view the backup history and logs, and click the row to expand its view.
3. Click the **Virtual Machines** tab.

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

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

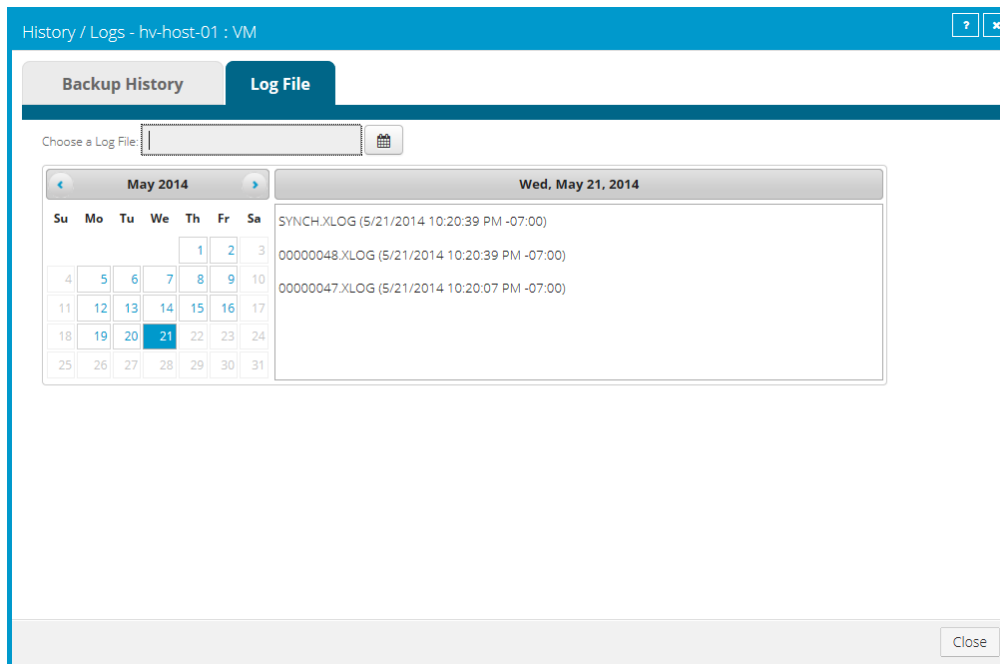
-  Not Protected — Indicates that the VM is not part of a backup job.
-  Never Run — Indicates that the VM is part of a backup job that was never run.
-  In Progress
-  Failed
-  Cancelled





4. Click the **Backup Status** column of the VM for which you want to view the backup history and logs.

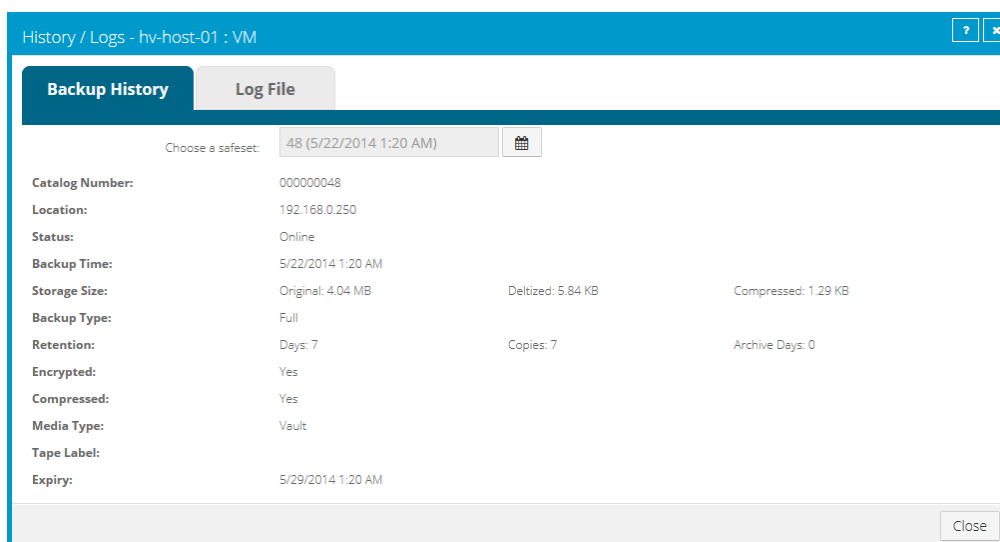


The **History / Logs** window lists log files from the date selected in the calendar.



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.



## 8.7 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. By default, the location is C:\Program Files\Carbonite Server Backup\Hyper-V Agent Management\Data\logs.

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 Carbonite.

## 9 Understanding and troubleshooting Hyper-V processes

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

### 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.

The Hyper-V Agent is optimized for backing up VMs whose virtual disks do not span multiple CSVs. Any VMs which span multiple CSVs will be backed up in a separate stage after all VMs which have their virtual disks on a single CSV.

### 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 contains mixed storage (e.g., one virtual disk on local storage and another virtual disk on a CSV). The Hyper-V Agent does not back up VMs that contain mixed storage. Other VMs in the same job are 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.

### Live migration fails

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

### VM restore fails

You cannot restore a VM that was backed up in a Windows 2012 R2 environment to a Windows 2012 environment.

## 10 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](#)

### 10.1 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>



**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.

To gather log files, click **File** menu and choose *Open log folder*. Compress the contents of the folder in a .zip file and attach it to your support request.

If the log archive and/or mail store 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 A: Alternate Hyper-V Agent deployments

### A1. 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 directly on a Hyper-V host, or on any supported Windows server that has local network access to the cluster. The server can be a physical or virtual machine, and must use the same DNS server as the Hyper-V cluster. Ideally, the server should be in the same active directory.

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.
- 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.

*Note:* 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.

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

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

### A2. Alternate deployment for protecting a Hyper-V standalone host

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

If you want to minimize performance impact in the environment, or you do not want to open the virtualized environment for Portal or vault access, you can install the Management service on a separate Windows server with local network access to the standalone host. The Management service server can be a physical or virtual machine that is on the same domain as the Hyper-V standalone host.

You must install the Host service on the standalone host.

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

*Note:* The Agent for Microsoft Windows cannot be installed on the standalone host.

## Appendix B: Install, upgrade and uninstall the Hyper-V Agent

This appendix includes procedures for installing and upgrading Hyper-V Agent components silently, and for uninstalling Hyper-V Agent components. For other installation information and procedures, see [Install and upgrade the Hyper-V Agent](#).

### B1. Install the Hyper-V Agent Management service in silent mode

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

```
installKitName /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
/L<localeID>	Optional. Specifies the language for installation log messages. The default value (1033 - English (United States)) is the only available value.  Note: Only English is supported with Hyper-V Agent 7.40. However, this version of the Agent can be installed on a non-English operating system. Support for other languages will be available in a subsequent release.
"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).
UIREG_NETADDRESS= <i>webUIAddress</i>	Specifies the host name or IP address of the Portal for managing the Hyper-V Agent.  Example: UIREG_NETADDRESS=192.0.2.233  Specifying the host name is recommended. This will allow DNS to handle IP address changes.

<code>UIREG_PORT=webUIportNumber</code>	<p>Optional. Specifies the port number used to communicate with Portal.</p> <p>Example: <code>UIREG_PORT=8086</code></p> <p>If you do not specify a <code>webUIportNumber</code>, port 8086 is used for communicating with Portal.</p>
<code>UIREG_USERNAME=webUIUser</code>	<p>Specifies the name of the Portal user associated with the Hyper-V Agent.</p> <p>Example: <code>UIREG_USERNAME=user@site.com</code></p>
<code>UIREG_PASSWORD=webUIUserPassword</code>	<p>Specifies the password of the specified Portal user.</p> <p>Example: <code>UIREG_PASSWORD=password1234</code></p>
<code>COORDINATOR_PORT=portNumber</code>	<p>Optional. Specifies the port used to communicate with Hyper-V Agent Host services.</p> <p>Example: <code>COORDINATOR_PORT=5444</code></p> <p>If you do not specify a port, port 5444 is used for communicating with Hyper-V Agent Host services.</p>
<code>INSTALLDIR=\"installFolder\"</code>	<p>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.</p> <p>Example: <code>INSTALLDIR=\"c:\Program Files\Management Service\"</code></p> <p>If you do not specify an installation folder, the Management service is installed in the default location (C:\Program Files\Carbonite Server Backup\...).</p>

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

```
Hyper-V_Agent_Management-x-xx-xxxx.exe /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\""
```

## B2. Install the Hyper-V Agent Host service in silent mode

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

```
installKitName /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&lt;localeID&gt;</code>	Optional. Specifies the language for installation log messages. The default value (1033 - English (United States)) is the only available value.  Note: Only English is supported with Hyper-V Agent 7.40. However, this version of the Agent can be installed on a non-English operating system. Support for other languages will be available in a subsequent release.
<code>\"<i>logFileName</i>"</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=<i>managementServiceAddress</i></code>	Specifies the host name 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 host name 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=<code>\"<i>installFolder</i>"</code></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=<code>\\"c:\Program Files\Host Service\"</code></code>  If you do not specify an installation folder, the Host service is installed in the default location (C:\Program Files\Carbonite Server Backup\...).

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

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

### B3. 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 /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 install the Management service in silent mode, you could run the following command:

```
Hyper-V_Agent_Management-x-xx-xxxx.exe /S /L1033 /V"/qn /L*v  
\"C:\logs\1.log\" "
```

### B4. 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 /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
\ " <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.

	<p>Example: <code>"C:\Logs\My Log.txt"</code></p> <p>If you do not specify a <code>logFileName</code>, the installation log is saved in the Windows installer default location (usually the user's temp directory).</p>
--	---

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

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

## B5. Uninstall the Hyper-V Agent Management service

*Note:* You can also uninstall the Management service using Add/Remove Programs in Windows Control Panel.

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. 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**.
3. On the **Welcome** page, click **Next**.
4. On the **Program Maintenance** page, select **Remove**, and then click **Next**.
5. On the **Remove the Program** page, click **Remove**.
6. If the **Files in Use** page appears, select **Automatically close and attempt to restart applications**, and then click **OK**.
7. On the **InstallShield Wizard Completed** page, click **Finish**.

## B6. Uninstall the Management service in silent mode

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

```
installKitName /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
-----------	-------------

Parameter	Description
<code>[/L&lt;localeID&gt;]</code>	<p>Optional. Specifies the language for uninstallation log messages. The default value (1033 - English (United States)) is the only available value.</p> <p><i>Note:</i> Only English is supported with Hyper-V Agent version 7.30. However, this version of the Agent can be installed on a non-English operating system. Support for other languages will be available in a subsequent release.</p>
<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>

## B7. 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.

*Note:* You can also uninstall the Host service using Add/Remove Programs in Windows Control Panel.

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. 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**.
3. On the **Welcome** page, click **Next**.
4. On the **Program Maintenance** page, click **Remove**, and then click **Next**.
5. On the **Remove the Program** page, click **Remove**.
6. If the **Files in Use** page appears, select **Automatically close and attempt to restart applications**, and then click **OK**.
7. On the **InstallShield Wizard Completed** page, click **Finish**.

## B8. 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 /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.

## Appendix C: 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 the Carbonite cloud.

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.

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.

### C1. 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.

*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 VMs 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.
5. Point to the VM that you want to find on the vault.  
A tooltip shows the name of the VM's task on the vault.

<input type="checkbox"/>	W2012	b-3	✓ Completed	today at 11:31 AM
<input type="checkbox"/>	zzz	VM ID: HV_MA6J2GNPC5CU5IIWQSAQ54BMM4	✓ Completed	today at 1:14 PM

6. To copy the name of the VMs' task on the vault, click the tooltip.  
The task name for the vault is copied to the clipboard.